

# Systemwide Commonalities in Market Liquidity

**Mark D. Flood**

Office of Financial Research

[Mark.Flood@treasury.gov](mailto:Mark.Flood@treasury.gov)

**John C. Liechty**

Office of Financial Research and Pennsylvania State University

[John.Liechty@treasury.gov](mailto:John.Liechty@treasury.gov)

**Thomas Piontek**

Office of Financial Research

[Thomas.Piontek@treasury.gov](mailto:Thomas.Piontek@treasury.gov)

The Office of Financial Research (OFR) Working Paper Series allows members of the OFR staff and their coauthors to disseminate preliminary research findings in a format intended to generate discussion and critical comments. Papers in the OFR Working Paper Series are works in progress and subject to revision.

Views and opinions expressed are those of the authors and do not necessarily represent official positions or policy of the OFR or Treasury. Comments and suggestions for improvements are welcome and should be directed to the authors. OFR working papers may be quoted without additional permission.

---

# Systemwide Commonalities in Market Liquidity

---

Mark D. Flood <sup>†</sup>

John C. Liechty <sup>‡</sup>

Thomas Piontek <sup>†</sup>

May 28, 2015

<sup>†</sup> – Office of Financial Research, Washington, DC

<sup>‡</sup> – Pennsylvania State U., State College, PA, and Office of Financial Research, Washington, DC

The authors thank Tobias Adrian, Greg Feldberg, Pete Kyle, Trish Mosser, Rob Engle, Tuomas Peltonen, Yves Schübert, and conference and workshop participants at Goethe University, Massachusetts Institute of Technology, Federal Reserve Bank of New York, Open Financial Data Group, the Office of Financial Research (OFR), the OFR–Federal Reserve Bank of Cleveland 2014 Financial Stability Conference, the New York University finance seminar, and the Institute for Pure and Applied Mathematics at the University of California at Los Angeles for numerous helpful comments and suggestions. Any remaining errors are the responsibility of the authors alone.

**Views and opinions expressed are those of the authors and do not necessarily represent official OFR or Treasury positions or policy, nor of any of the authors' affiliates or employers.**

# Systemwide Commonalities in Market Liquidity

## Abstract

We explore statistical commonalities among granular measures of market liquidity with the goal of illuminating systemwide patterns in aggregate liquidity. We calculate daily invariant price impacts described by Kyle and Obizhaeva [2014] to assemble a granular panel of liquidity measures for equity, corporate bond, and futures markets. We estimate Bayesian models of hidden Markov chains and use Markov chain Monte Carlo analysis to measure the latent structure governing liquidity at the systemwide level. Three latent liquidity regimes — high, medium, and low price-impact — are adequate to describe each of the markets. Focusing on the equities subpanel, we test whether a collection of systemwide market summaries can recover the estimated liquidity dynamics. This allows an economically meaningful attribution of the latent liquidity states and yields meaningful predictions of liquidity disruptions as far as 15 trading days in advance of the 2008 financial crisis.

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
1.1	The Challenges of Liquidity Measurement . . . . .	5
1.2	Liquidity Measurement in Practice . . . . .	13
<b>2</b>	<b>Model Description</b>	<b>20</b>
2.1	Market Liquidity Measures . . . . .	20
2.2	Univariate Models of Latent Structure . . . . .	30
2.3	Hierarchical Model . . . . .	38
<b>3</b>	<b>Data</b>	<b>40</b>
<b>4</b>	<b>Liquidity Regimes</b>	<b>43</b>
4.1	Individual Market Liquidity . . . . .	43
4.2	Explaining Liquidity Regimes . . . . .	50
<b>5</b>	<b>Predicting Liquidity Regimes</b>	<b>58</b>
<b>6</b>	<b>Conclusion</b>	<b>63</b>
	<b>References</b>	<b>65</b>

# 1 Introduction

Liquidity is a central consideration for market quality in general, and for financial stability in particular. We present a new approach to the study of liquidity that identifies broad-based patterns in daily data for individual markets to help explain aggregate, systemwide liquidity conditions. Although funding liquidity in the wholesale markets for institutions is the most immediate concern for systemwide conditions Brunnermeier and Pedersen [2009], there are many more individual markets for financial assets than for intermediaries' liabilities. It is an empirical question whether there is additional information in the much larger panel of asset markets that might help to explain liquidity conditions in the funding markets. We expand on previous studies of commonalities in liquidity — e.g., Chordia et al. [2000] and Karolyi et al. [2012], who find that there are indeed significant patterns in the detailed data — by analyzing liquidity in a range of asset classes, including equities, bonds, and financial futures.<sup>1</sup> We also extend the commonalities approach with a novel methodology for connecting aggregate liquidity patterns to a panel of systemwide market summaries.

Starting with granular measures of market liquidity, based on the recent invariant price-impact measures of Kyle and Obizhaeva [2014], we estimate a daily panel of liquidity conditions across a broad range of markets. Specifically, our initial implementation considers volatility index futures, oil futures, and sector portfolios for the Center for Research in Securities Prices (CRSP) universe of U.S. equities and the Transaction Reporting and Compliance Engine (TRACE) universe of corporate bonds over the decade 2004-14. The market-invariant approach of Kyle and Obizhaeva [2014] carefully normalizes for local volume and volatility conditions to produce liquidity measures that are directly comparable

---

<sup>1</sup>Most prior studies of commonality have focused on equities markets alone. An exception is the recent working paper by Marra [2013] which pairs individual equities with their matching credit default swaps (CDS); her emphasis, however, is on firm-level interactions between the securities rather than systemwide liquidity.

across markets and order-flow conditions. Comparability is crucial for aggregating local liquidity conditions to support systemwide analysis. Using this panel of daily liquidity measurements, we estimate Bayesian hidden Markov chains (HMC) models, using Markov chain Monte Carlo (MCMC) inference methods, to capture the latent structure of each series, and assess the latent structure governing liquidity at the systemwide level. The HMC approach posits that the dynamics of each daily price impact measure (33 in our sample) are determined by an underlying variable that alternates among several liquidity states to drive sudden changes in the observed levels of price impact. The underlying states are latent — i.e., not directly observable — and must be inferred from the dynamics of daily price-impact measurements. In the initial analysis, we estimate each price impact series independently; that is, we assume no coordination between the dynamics of the latent liquidity states across markets. Nonetheless, we find surprising consistency in the dynamics of market liquidity across all of these markets. From the perspective of a policy maker who seeks to identify, or even predict, turbulent episodes in the financial system, we find that three liquidity regimes are adequate to describe each market: high, intermediate, and low. Moreover, we find that the low liquidity regime afflicts all markets roughly simultaneously during the financial crisis of 2008.

Studies of market liquidity are typically grounded in practical considerations about market quality and exploit local microstructural characteristics to craft market-specific liquidity measures.<sup>2</sup> This tendency towards customization of liquidity metrics supports the economic interpretation of results and makes full use of the available data. It is a hindrance for a systemwide analysis, however, because data availability and the microstructural interpretation of the estimated measures can vary considerably from one market to the next. Nonetheless, as a point of comparison, we also consider the dynamics of four tra-

---

<sup>2</sup>The literature on market liquidity is too extensive for a fully comprehensive survey. Recent overviews include Foucault et al. [2013], Holden et al. [2014], Gabrielsen et al. [2011], Hasbrouck [2007], and Stoll [2000] in his American Finance Association Presidential Address. In addition, Amihud et al. [2013] compile a set of classic papers.

ditional measures of liquidity: the implied bid-ask spread of Roll [1984]; Kyle’s- $\lambda$  (see Kyle [1985]); the volume-scaled absolute return measure of Amihud [2002]; and log volatility.<sup>3</sup> While these alternative measures indeed capture certain patterns in liquidity dynamics, they also exhibit clear statistical anomalies and occasional structural changes related to microstructural innovations. Among the current class of market liquidity measures available in the literature, the market-invariant measure most closely meets the requirement of comparability over time and across markets.

We build on earlier studies of liquidity commonalities by linking the estimated latent liquidity states from multiple markets together with a multinomial probit model driven by a collection of daily summary (i.e., system-level) series. This phase of the analysis restricts attention to the equities markets to avoid complications arising from missing values.<sup>4</sup> This framework permits an assessment of summary series as potential advance indicators of systemic illiquidity. The model reveals that a number of summary series, including the Dow Jones U.S. Real Estate Industry Group Index, Treasury-Eurodollar (TED) spread, VIX<sup>®</sup>, and the S&P 500 price-to-book (P/B) ratio are statistically significant in explaining liquidity in these equity markets. These summary series are observed daily and unsurprisingly exhibit a high level of temporal correlation; as a result, adding lagged values does not improve the performance of the probit model. On the other hand, the same high level of temporal correlation allows a version of the model based on lagged summary series to predict sudden future shifts in the liquidity states — as much as 15 days in advance of the liquidity crisis in September 2008. Although a detailed exploration of the predictive power of these methods is a point of future research, these preliminary results suggest the method might support market monitoring and early warning systems for illiquidity episodes.

---

<sup>3</sup>The NYU Volatility Institute (V-Lab) [2014] at New York University provides calculated time series for three variants of the Amihud [2002] measure. We include simple log volatility in the comparison because of the prominent role of volatility in the Kyle and Obizhaeva [2014] measure.

<sup>4</sup>We are exploring techniques to address the problem of missing observations; this remains a topic for future research.

The paper is structured as follows. The remainder of Section 1 discusses the general issue of liquidity measurement, and provides a rationale for our approach. Section 2, describes the models and sampling strategies for the MCMC analysis for both the univariate models (i.e., one market at a time) and the hierarchical model (multiple markets at a time). Section 3 describes the data and specific formulas for measuring price impact. Section 4 reports the findings from aggregating both the market specific analyses and the multiple-market analysis. Section 5 discusses some alternative modeling approaches (e.g., vector autoregressive models) and their limitations with respect to policy applications, but also touches on how these might fit into a broader modeling approach that focuses on prediction of liquidity dynamics; this discussion ends with a demonstration of the predictive power of the probit model during the crisis of 2008. We conclude with a discussion of potential future work, including ways that limitations of the current approach might be overcome to improve our ability to forecast liquidity.

## 1.1 The Challenges of Liquidity Measurement

Conceptually, liquidity is the ease with which participants in the financial system can convert their claims to cash. The settlement obligations of the vast majority of financial contracts are stated in terms of cash: payors must deliver cash, and payees must accept it. As a result, the ability of market participants to “get to cash” has important implications for the overall functioning of the system. Liquidity is particularly important for the study of financial stability, as sudden shifts in liquidity have historically been a defining characteristics of financial crises.<sup>5</sup> We focus on the interplay between two key aspects of

---

<sup>5</sup>Kindleberger [1993, ch. 15], for example, recounts the history of crises in Western economies, with a special focus on the lender of last resort as a provider of backstop liquidity to the system. Reinhart and Rogoff [2009], and Schularick and Taylor [2012] consider a similar sample, with the latter focusing on the role of credit booms gone wrong as a precursor to crisis events. Because aggregate credit growth is typically facilitated by expanding bank balance sheets, there is an important empirical connection linking credit cycles, leverage cycles, and liquidity cycles.

systemwide liquidity: conditions at the aggregate level, including wholesale funding markets; and the information available on liquidity in the much more numerous individual asset markets.

Market liquidity arises because there are agents who stand ready to buy (or sell) an asset. Arbitrage implies that there should be profits available to those with robust valuation models and accurate information to feed them, so counterparties in a liquid market should be easy to find for a modest price concession. In principle, liquidity providers should step in opportunistically if the asset is offered at an acceptable discount, but will have little incentive to pay more than the current market price. The practical upshot for liquidity measurement is that we measure the extent of illiquidity in the system as one-sided deviations from an ideal benchmark of perfect liquidity. It is important to consider what constitutes a reasonable price in this context. Most financial instruments have limited liability, so there is typically some nonnegative price at which buyers come forward; however, a market is not liquid if buyers only appear for fire-sale offers. Financial engineers can often provide plausible mark-to-model valuations, even for contracts that trade infrequently.<sup>6</sup>

Financial institutions aggregate and reallocate liquidity — i.e., the cash available to and from liquidity providers in local markets. Systemic imbalances in liquidity therefore tend to appear in wholesale funding markets, where they are a commonplace feature of financial crises. Allen and Gale [2009] distinguish the role in crises of liquidity fundamentals (e.g., subprime mortgage valuations), which operate primarily in asset markets, and panics (e.g., bank runs), which primarily affect funding markets. Both forces might be present in any particular episode, and measuring their influence is an important empirical task. Ideally, measures of liquidity to support financial stability monitoring would be both timely — available at high frequency to track developments in near real time — and forward-looking

---

<sup>6</sup>The recent Presidential Address to the Econometric Society by Holmström [2012] emphasizes the lengths markets will endure to achieve liquidity by forcing assets to be “informationally insensitive.” See also Dang et al. [2012].



— possessing some forecasting power to serve as an early warning signal. These goals are often defeated in practice by certain fundamental challenges. In particular, liquidity exhibits three interrelated characteristics that present special complications to measurement: latency, nonlinearity, and endogeneity. Each of these challenges has ramifications for both funding liquidity and market liquidity.

### **1.1.1 Latency**

Latency means that much of the most interesting liquidity behavior is ex-ante unobservable. At the microstructural level, we typically most wish to know not merely the prices of recent trades or the current best bid and offer, but how deep or resilient the market will be in the presence of unusual order flow. This implies that the most interesting liquidity events are typically also the rarest. Moreover, many trading mechanisms encourage the participation of liquidity providers — buyers and sellers — by restricting information availability. Where the microstructure relies on limit orders, techniques to limit transparency include closed limit-order books and hidden (or “iceberg”) orders; e.g., [Parlour and Seppi, 2008, Section 2.6], Bessembinder et al. [2009]. More generally, markets can utilize anonymous brokerage to conceal trader identities, and/or limited-access upstairs trading venues for large trades; e.g., Degryse et al. [2014], Nimalendran and Ray [2014], Zhu [2013], and Foley et al. [2013].

Moreover, traders are not compelled to reveal their intentions by actually issuing (or canceling) a limit order prior to the moment of truth; some version of a market order is typically available. Conversely, private visibility into their own customers’ positions and order flow can be a valuable information source for dealers as they try to avoid adverse selection by those they trade with; see, for example, Kyle [1985], Easley et al. [1996], and Evans and Lyons [2002]. The inability of researchers to measure these aspects of liquidity directly makes these forces a natural subject for empirical modeling. If key drivers, such as dealer intentions or customer order flow, are not immediately observable, latent structure

might be recovered through statistical inference.<sup>7</sup>

To hedge against liquidity surprises, firms frequently arrange for contingent liquidity in the form of lines of credit or derivative contracts. However, hedgers must still worry about the wrong-way risk that their supposed guarantors will themselves fail under the precisely the event being insured against. Liquidity is therefore never purely localized to a single transaction or market. At the broadest levels, global liquidity depends in part on the responses of firms (and policy makers) to situations they may have not have planned for explicitly. For example, it is difficult to know in advance whether an initial deleveraging event will generate enough selling pressure to create a fire-sale feedback loop. Geanakoplos [2003], for example, emphasizes that collateral margins are likely to bind in a crisis, unexpectedly depriving participants of flexibility at the crucial moment. Similarly, the repeated clearing crises of the nineteenth century (see Calomiris and Gorton [1991]) were invisible to bankers in the system until it was too late. In general, liquidity measurement is “more honored in the breach,” in the sense that it is easier to assess the depth or resilience of the market when conditions are stressful enough to violate the perfectly liquid ideal.

We address the challenge of latency by adopting a MCMC approach to estimate the latent regime structure governing the observed price impact series. The maintained assumption is that, while the markets’ liquidity behavior is indeed largely latent, these hidden patterns will reveal themselves in a broad cross-section of markets observed at relatively high frequency (daily). In the results below, we are indeed able to identify three meaningful latent liquidity states (high, medium, and low price-impact) that seem to govern the observed liquidity behavior.

---

<sup>7</sup>Examples include: in equities markets, Hasbrouck and Saar [2009]; in corporate bond markets, Mahanti et al. [2008]; and in interbank funding markets, Gefang et al. [2011].

### 1.1.2 Nonlinearity

Nonlinearity in the response of liquidity to significant market changes compounds the problem of unobserved behavior. Nonlinearity is a challenge for liquidity measurement because it hampers our ability to extrapolate from small-scale, localized effects to the larger, out-of-sample effects that are often of greatest concern. Numerous studies have documented the empirical regularity that price response to order flow tends to be concave function of the transaction size. Intuitively, order flow can move the price significantly before additional liquidity providers arrive to dampen the effect. Much of the literature identifies a square-root rule that posits price impact to be proportional to the square root of the transaction size.<sup>8</sup> In contrast, recent work by Kyle and Obizhaeva [2011a,b] argues for a cube-root rule.

Most of the literature goes beyond describing these basic reduced-form empirical regularities to identify specific underlying behaviors or mechanisms that could be at work. For example, limit orders (and other contingent liquidity) may crowd behind the best posted quote, so that an order flow impulse large enough to work through this initial phalanx would expose gaps in the book, provoking an abrupt shift in prices. Such unevenness in market depth may be an important source of fat-tailed returns distributions. Kyle and Obizhaeva [2014] provides a theoretical justification for their cube-root rule, based on Poisson arrivals of speculative order flow. A recent paper by Bookstaber et al. [2015] points to asymmetries in decision-response times between buyers and sellers as a possible source of nonlinearities in price impact. As a possible example of this mechanism, they point to the 1987 market crash, in which relatively speedy portfolio insurance traders in the index futures markets overwhelmed the order-flow capacity of traditional equities dealers as program traders laid off the inventory in the spot market.<sup>9</sup> Duffie [2010] suggests three

---

<sup>8</sup>See for example, Gabaix et al. [2006], Bouchaud et al. [2008], Hasbrouck and Seppi [2001], and Toth et al. [2011].

<sup>9</sup>The empirical fact that the bulk of price discovery for traded equity indexes occurs in the futures, not

categories of diverse explanations for the various delays across markets in the arrival of a countermanding response to an initial order-flow impulse: search; dealer capital constraints; and investor inattention. In extreme cases, a large initial price move may repel, rather than attract, price-stabilizing speculative order flow; DeLong et al. [1990] present a key early model of such positive-feedback trading. Models in this tradition are similar in spirit to the “momentum” trading explanation of Jegadeesh and Titman [1993, 2001], in that the driving force for current trading behavior is the recent history of prices alone.

Nonlinearity in liquidity is an even greater worry for systemic stability, where the stakes are correspondingly higher. At this level, interactions among nodes in the system can conspire to produce self-amplifying feedback loops. Tirole [2011] provides a tour of systemic pathologies related to illiquidity, including contagion, fire sales, and market freezes. He underscores the central fact that one of the basic services provided by the banking (and shadow banking) sector — namely maturity transformation — render it especially vulnerable to runs and other liquidity surprises. Brunnermeier [2009] provides a good overview of how these forces played out in practice, at least through the early (and most severe) phases of the recent crisis. He highlights four specific channels: (a) deleveraging spirals driven by erosion in capital levels and increases in lending standards and margins; (b) a credit crunch motivated by funding concerns; (c) outright runs, exemplified by Bear Stearns and Lehman Brothers; and (d) netting failures due to real or perceived counterparty credit risks. All these modalities involve liquidity. Adrian and Shin [2010] emphasize the role played by institutional leverage in both the expansion and contraction of the system. Figure 1, adapted from similar illustrations in Adrian et al. [2013a], illustrates clearly that institutional leverage expands and contracts via adjustments to assets and liabilities—not equity — suggesting that increases in leverage correspond to increases in overall liquidity, since bank deposits and other liabilities are a key component of liquid assets for other

---

the spot markets, has long been recognized. See, for example, Hasbrouck [2003]. The seminal paper in this vein is Kawaller et al. [1987].

participants in the system. Leverage is strongly procyclical; by increasing global liquidity, aggregate balance sheet expansion encourages investment spending and tends to ease margin constraints. Figure 1 also indicates a correlation between bank leverage and market volatility, at least through the course of the most recent business cycle: periods of low volatility (green and yellow markers) tend to correspond to increases in bank leverage, and episodes of high volatility (orange and red) tend to match decreases in leverage. Unfortunately, the volatility-leverage-liquidity spiral works in reverse as well, leading to debt overhangs as the system contracts, with associated increases in institutional risk aversion and liquidity hoarding.

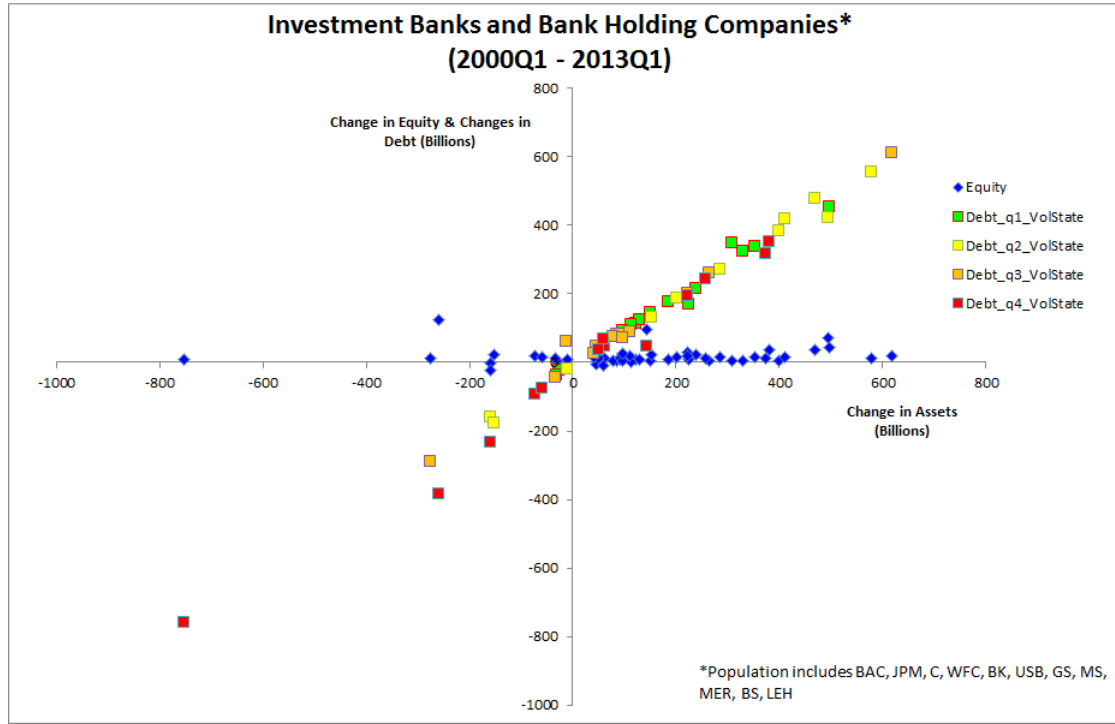


Figure 1: The Leverage-Liquidity Cycle

Sources: Bloomberg L.P., OFR analysis

We address the challenge of nonlinearity by agnostically allowing the data determine the correct number of liquidity states for each time series. Notably, for all 33 of our univariate series, three liquidity states are adequate to explain the observed variation in the price-impact statistics. Because the expected price impact is allowed to vary idiosyncratically

for each of the three regimes (high, medium, and low price-impact), this model naturally captures nonlinearities in price impact.

### 1.1.3 Endogeneity

Endogeneity means that liquidity is partly a network effect that emerges through the interactions of many market participants. Thus, active markets should have less need for identified liquidity providers who convert investments to cash by purchasing or rediscounting others' financial assets. Endogenous liquidity creates a straightforward network externality in the sense of Pagano [1989] and Economides [1996]: investors are naturally more willing to enter markets where there are already many other traders and large transaction volumes, because this provides an implicit assurance that counterparties will be easy to find when needed. A familiar example of this phenomenon is the contrast between trading for on-the-run and off-the-run Treasuries; see Barclay et al. [2006]. Similarly, Bessembinder et al. [2006] find that liquidity externalities are consistent with the significantly reduced trade execution costs that followed the introduction of the TRACE feed, which increased transparency in the corporate bond market. Liquidity externalities are also often touted as a benefit of high-frequency trading.

Liquidity externalities operate at the level of the system as well.<sup>10</sup> They have long been a central concern of financial stability supervisors. Elliott et al. [2013], for example, document this multi-dimensional U.S. regulatory history. One of the core functions of central banking is to provide a lender of last resort in a crisis, an idea first flirted with as an expedient in London's panic of 1825, codified in Bagehot [1873], and institutionalized in the United States with the creation of the Federal Reserve in 1913. Recourse to a potentially unlimited source of liquid cash from outside the network of banks and financial firms is important in a crisis, when information asymmetries and other constraints prevent firms from liquidating

---

<sup>10</sup>There are many discussion of this endogenous systemic externality. See, for example, Morris and Shin [2004], Dang et al. [2010], and Adrian and Shin [2010] and the references therein.

their assets to meet withdrawals. On the other hand, some have argued that this explicit promise of effectively unlimited contingent liquidity creates a moral hazard — that too-big-to-fail banks undertake excessive leverage and maturity transformation, comfortable that the Fed’s emergency backstop provides them with a free liquidity put.<sup>11</sup> Moreover, in spite of discount window access, Cornett et al. [2011] find that banks in the recent crisis were forced onto a more defensive liquidity posture, in part because Lehman Brothers’ failure diverted commercial paper borrowers to draw on banks’ liquid assets via backup lines of credit, and partly because wholesale funding sources suddenly shrank. A net result was a restriction in commercial lending. In the wake of the crisis, macroprudential supervisors have focused renewed attention on liquidity buffers, including the new Basel III requirements for banks to maintain net stable funding and liquidity coverage ratios.<sup>12</sup>

We address endogeneity by estimating a hierarchical model that searches for common liquidity structure throughout the cross-section of observed price-impact series. Although this work is in preliminary stages, and is currently limited to the cross-section of equity markets, we are able to identify significant patterns and attribute them statistically to particular systemwide market summaries, providing some economic interpretation for the estimation.

## 1.2 Liquidity Measurement in Practice

How one measures liquidity depends in part on the portion of the financial system under examination. As noted, much of the literature focuses either on liquidity in the narrow

---

<sup>11</sup>Goodhart [2008] and Farhi and Tirole [2009] have made this argument. The term “liquidity put” is a metaphor for a commonly used recourse covenant that allows investors in a partially debt-funded structured investment vehicle (SIV) to put back their shares in the SIV to the sponsoring bank if the SIV is unable to roll over its short-term debt; see, for example, Entwistle and Beemer [2008]. However, during at least one episode — the Y2K millennium date change — the Fed literally sold liquidity put options; see Sundaresan and Wang [2006].

<sup>12</sup>The new liquidity framework for banks is discussed in Basel Committee on Banking Supervision [2013, 2010], Adrian et al. [2013b], and Bank of England [2011].

context of a particular financial market — so-called market liquidity — or at the aggregate level of the financial system as a whole — so-called global or funding liquidity. Brunnermeier and Pedersen [2009] connect these two strands of the literature, using bank balance sheets as an organizing device, as depicted in Figure 2. In this framework, the distinction between market and funding liquidity is based essentially on which side of an intermediary’s balance sheet is involved. Market liquidity refers to the ease with which financial institutions can convert securities or loans from their asset portfolio to cash.<sup>13</sup> Financial institutions can participate as both suppliers and demanders of liquidity in these markets, as indicated by the bidirectional arrows on the left side of Figure 2.

Funding liquidity refers to the ease with which institutions can obtain cash by borrowing in funding markets. Figure 2 underscores that one of the central functions of banks and similar intermediaries is to convert relatively long-maturity, low-liquidity commitments on the asset side to relatively short-maturity, high-liquidity obligations on the liability side of intermediaries’ balance sheets. Official liquidity represents the range of short-term cash resources available in a financial crisis — when the wholesale funding markets fail — from central banks and other agencies, enterprises, and programs with explicit or implicit taxpayer backing.<sup>14</sup> Because these liquidity resources typically come into play only in unusual but important occasions, they appear as dashed arrows in the figure, which flow in only one direction.

---

<sup>13</sup>Harking back to Moulton [1918], Mehrling [2010] refers to this sort of liquidity as the “shiftability” of bank or dealer assets — i.e., the ability to shift them into cash. The asset side of Figure 2 represents a primary point of contact between the financial system and the real economy. Commitments like corporate or mortgage loans typically translate directly into real activity such as workforce expansions and home improvements investments. Cornett et al. [2011] analyze market and funding liquidity empirically, along with their net effect on overall credit supply.

<sup>14</sup>In the United States, official or “outside” liquidity includes deposit insurance, the Fed’s discount window, Federal Home Loan Bank advances, as well as the numerous emergency facilities created as expedients in the recent crisis; see Fleming [2012] for details. For a definition and model of inside and outside liquidity, see Holmström and Tirole [2013].



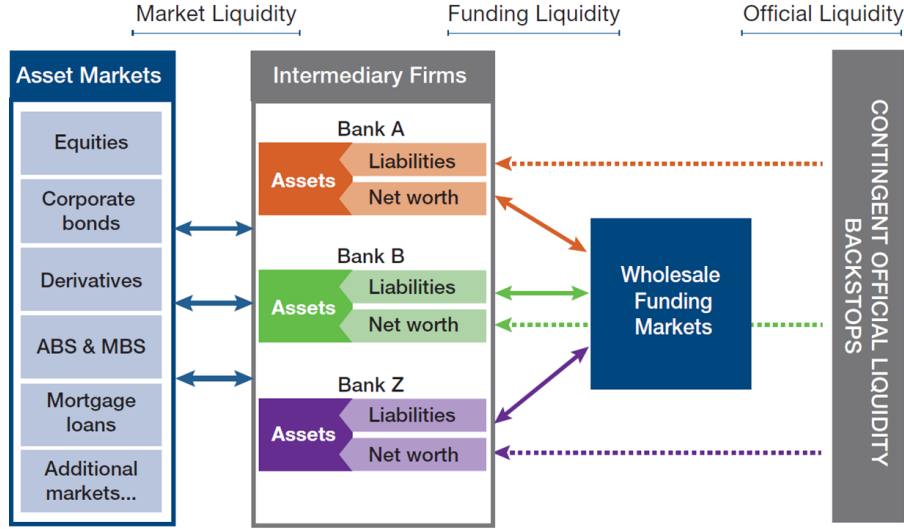


Figure 2: Liquidity Transformation via Financial Intermediation  
Source: OFR analysis

### 1.2.1 Aggregate Liquidity Measures

An institution typically handles its market liquidity needs by drawing on its own cash reserves or selling assets for cash. If the ordinary give and take of trading in the asset markets does not net out, the firm can turn to the funding markets to borrow or lend the difference. Wholesale funding markets thus aggregate much of the endogenous net supply and demand of liquidity overall, and prices in these markets provide a bellwether for the state of system. Figure 3 depicts several commonly used measures of aggregate liquidity conditions derived from prices in wholesale funding markets. Following Brunnermeier [2009], Boudt et al. [2013], and Boyson et al. [2010], we proxy the TED spread as the difference between three-month T-bill yields and three-month LIBOR.<sup>15</sup> A frequently cited alternate spread measure for funding liquidity conditions is the LIBOR-OIS (London interbank offered rate vs. overnight index swap) spread; see Gefang et al. [2011], Michaud and Upper [2008], and Taylor and Williams [2009]. Both spreads capture deviations of

<sup>15</sup>Futures traders at the Chicago Mercantile Exchange first noted the TED spread in the early 1980s, where they tracked the pricing differential between the three-month T-bill futures and three-month Eurodollar futures contracts, which traded in neighboring pits McCauley [2001].

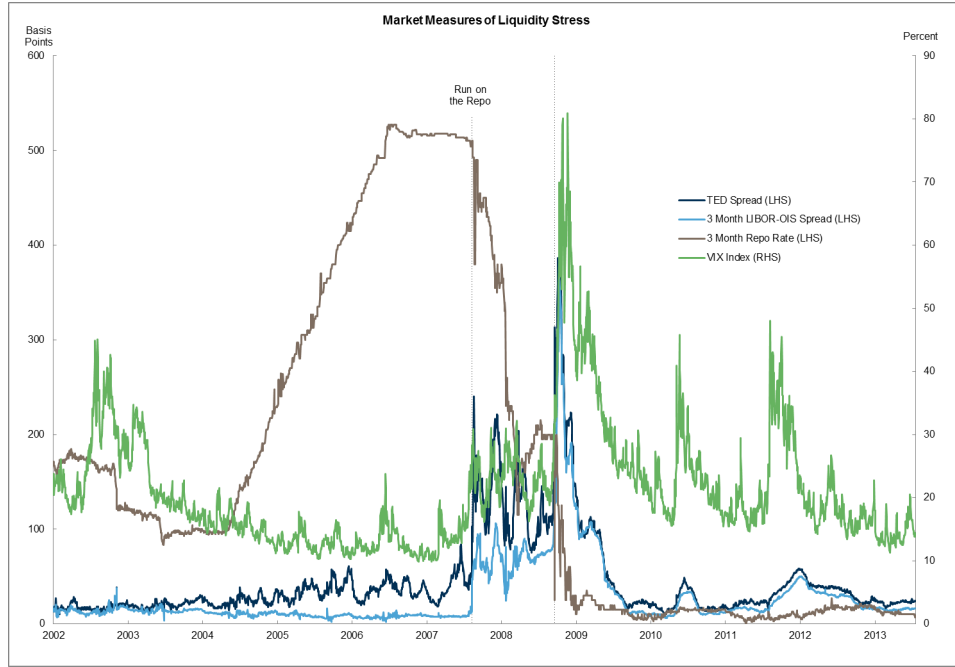


Figure 3: Bellwether Liquidity Measures in Wholesale Funding Markets  
Source: Bloomberg L.P.

borrowing conditions in the interbank markets from “pure” credit-risk-free borrowing, reflecting reticence to lend to banks for various reasons, including credit risk and aggregate liquidity anomalies. The VIX<sup>®</sup> is a traded index of market volatility. Its gradual downward drift during the pre-crisis period (roughly 2002 to 2007) is symptomatic of the so-called volatility paradox: market risk as measured by price volatility was dropping while overall risk exposures were simultaneously (and not coincidentally) building across the system (Brunnermeier and Sannikov [2012]).

The first major foreshock of the crisis came in August 2007, triggered by an absence of liquidity that prevented BNP Paribas from marking to market several of its investment funds backed by subprime mortgages. The surprise provoked a sharp but very temporary drop in the three-month repo rate and a simultaneous more permanent jump in both the TED and LIBOR-OIS spreads. Market turmoil persisted through the failure of Bear Stearns in March 2008, until the failure of Lehman Brothers in September 2008 raises

liquidity problems to a new level. Interest rate spreads provided a high-frequency glimpse into aggregate liquidity conditions, but price moves only hinted at the underlying changes in cash holdings. Engle et al. [2012] estimated a multiplicative error model (MEM) to disentangle the interrelated forces of liquidity (order-book depth) and volatility in the Treasuries market. By this measure, liquidity dropped sharply and volatility spiked during the crisis.

As financial firms withdrew, the aggregate endogenous liquidity supply in the wholesale funding markets became inadequate to satisfy current cash obligations, and official liquidity providers were forced to step in. The Federal Home Loan Banks (FHLBs), saw early increases in advance funding to member institutions, which include most large commercial banks (Ashcraft et al. [2010]). Over the first few months after August 2007, FHLB advance funding increased by more than \$200 billion. Although this was modest compared to what was to come, it was a significant departure from business as usual at the time. The full-blown crisis emerged with the Lehman failure in September 2008. Wholesale markets collapsed, and financial firms proceeded en masse to the Fed’s backstop liquidity programs.<sup>16</sup> Surprisingly, the Fed’s signature lender-of-last-resort facility, the discount window, played a miniscule role throughout the episode.

The sharp shift in supply and demand for wholesale liquidity also showed up in the management of banks’ cash reserves, as depicted in Figure 4. Prior to the Lehman shock, aggregate bank reserve balances (which paid no interest prior to July 2013) hovered near zero; banks continued to rely on wholesale funding markets to meet short-run cash contingencies. After the Lehman failure, banks begin to hold precautionary reserve balances

---

<sup>16</sup>The Bear Stearns failure necessitated recourse to the Fed, including the creation of two brand new liquidity vehicles, the Primary Dealer Credit Facility (PDCF) and the Term Securities Lending Facility (TSLF). Following the Lehman failure, the brunt of the wave of new funding demands was borne initially by the Fed’s swaps facility, along with FHLB advances, the PDCF and TSLF again, and the newly created ABCP MMMF liquidity facility (AMLF). Much of this funding subsequently transitioned to other programs, including the Treasury’s new Troubled Asset Relief Program (TARP), created in October 2008. For an analysis of the Fed’s various large-scale asset purchase programs, see Chen et al. [2012a], D’Amico and King [2012], and D’Amico et al. [2012].

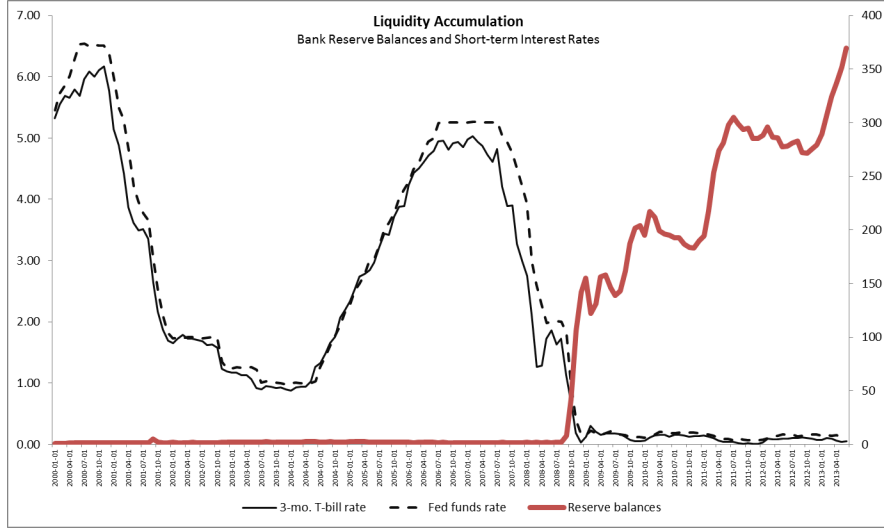


Figure 4: Bank Reserve Balances and Short-term Interest Rates  
Source: St. Louis Federal Reserve Economic Data

in significant quantities; this practice of reserving has continued essentially unabated. At the same time, the Federal Reserve has flooded markets with liquidity, driving yields on overnight Fed funds and T-bills to near zero. The persistence of the short-term riskless rate near the zero lower bound while loanable funds pool up — potential lenders always have the alternative of holding cash instead of accepting a negative return — suggests a market failure.

### 1.2.2 Granular Liquidity Measurement

At the aggregate level of funding liquidity, a primary concern is whether the financial system has the internal flexibility to satisfy all of its immediate funding needs. By design, liquidity measures based on prices and volumes in the wholesale funding markets aggregate information from across the financial system. There are only a relative handful of funding products traded in these markets, which are dominated by a small set of very large institutions. Economic equilibrium means that every borrower in the wholesale funding markets should be able to find a willing lender, but the long and painful history of systemic crises demonstrates that this equilibrium is not reliable. This is the purview of central banking,

macroprudential regulation, and systemic supervision.

Yet liquidity also applies to the many thousands of markets for equities, bonds, indexes, commodities, and derivatives. Market liquidity focuses on the intricacies of these (individually) smaller markets. Although these smaller markets are not as immediately connected to system-level stability, they are much more numerous than the interbank funding markets. The availability of granular, high-frequency data on transaction prices, bid-ask quotes, trading volumes and customer order flows facilitates detailed modeling of the behavior of market participants. Liquidity metrics for these markets may therefore carry additional information about overall liquidity that is lost in the aggregation to the wholesale level. In particular, asset markets offer a smorgasbord of different industries, product types, geographic concentrations, maturity habitats and credit grades that is not available in the short-term, interbank funding markets. It is therefore an empirical question whether this diversity generates measurable cross-sectional patterns in liquidity, and whether this cross-sectional information is helpful in understanding systemic behavior.

Our work builds on earlier studies that look for aggregate liquidity patterns. Chordia et al. [2000] was the first in a series of papers to search for “commonalities in liquidity” in the cross-section of equity markets. They perform time-series regressions of liquidity, measured as market depth and bid-ask spreads, for individual stocks on cross-sectional average measures of liquidity. The data is noisy —  $R^2$ s are low — but there is strong evidence of contemporaneous correlation between individual stocks and the aggregate. Karolyi et al. [2012] extend the analysis to an international comparison of thousands of stocks in 40 countries. Again, commonalities in liquidity exist, and unsurprisingly differ significantly across countries and over time.<sup>17</sup> Recent research by Chen et al. [2012b] combines price

---

<sup>17</sup>Karolyi et al. [2012] define commonality by the  $R^2$  of each stock’s daily price-impact measure, per Amihud [2002], on the average price impact for all other stocks in the country. Individual stock commonality measures are averaged to get a country-level commonality index. Karolyi et al. [2012, p. 99] attribute the time-series variation to both supply- and demand-side proxies in funding markets via regression analysis, noting that “demand-side explanations are more reliably significant.”

information from financial markets with quarterly quantity information from the Flow of Funds data in an effort to distinguish the differential impact of shifts in liquidity demand versus liquidity supply. They distinguish between core and noncore liquidity, where the noncore category includes financial firms’ liabilities held by other financial institutions.

## 2 Model Description

This section describes our choice of a market liquidity metric, how we apply MCMC analysis to estimate implicit liquidity states, and how we aggregate information across markets to detect systemwide patterns.

### 2.1 Market Liquidity Measures

A central goal of this research is to identify broad patterns or commonalities in market liquidity that might support a formal program for monitoring systemwide liquidity conditions. This implies a difference in scope from earlier studies of liquidity commonality. By casting a wide net across diverse instrument types, we hope to have a better chance of detecting emerging liquidity anomalies and identifying key liquidity indicators and important patterns among the markets being monitored. It is impossible to know with certainty *ex ante* which market(s) might participate in a salient way in a systemic illiquidity event. Therefore, the cross-section of asset markets should be both broad and extensible. To be responsive to evolving liquidity conditions as a systemic surveillance tool, the liquidity measure should be available at (at least) a daily frequency. These considerations translate into four minimal requirements for an acceptable market liquidity metric for our purposes:

- Feasibility – The data inputs needed to calculate the metric should be available for a broad range of markets.

- Timeliness – It should be practical to update the metric with at least a daily frequency for all markets in the sample.
- Comparability – The metric should have the same general statistical characteristics (e.g., scale and dimension) for all markets to which it is applied to support comparisons and aggregation across a broad range of markets.
- Granularity – The measurements should be resolvable to the level of individual markets, to support attribution of systemic liquidity events to specific sectors or markets.

These criteria narrow the field of candidate measures from the research literature considerably.<sup>18</sup> As a practical matter, the feasibility requirement restricts attention to those metrics that depend only on prices and volumes (or derived values, such as returns and volatility), because most markets have post-trade transparency of this information. Metrics requiring pre-trade transparency (e.g., quoted bid-ask spreads), customer order flow, or dealer inventories are not feasible by this definition.<sup>19</sup> Conditional on satisfying the feasibility requirement, the timeliness requirement usually does not bind. Even metrics requiring a multi-day estimation interval, such as the regression model for Kyle’s lambda, can employ a rolling window to produce daily liquidity observations.

The comparability requirement is important for a systemwide analysis. Most metrics envision measuring liquidity in one market at a time (thus providing granularity). In contrast, systemic monitoring requires the ability to understand not only whether illiquidity in a given market is unusual relative to its *own* history, but also relative to conditions in *other* markets. For many metrics, comparability might be achieved by an appropriate market-specific normalization, but other metrics are more problematic in this regard.

---

<sup>18</sup>The overviews by Foucault et al. [2013], Holden et al. [2014], and Gabrielsen et al. [2011] provide a universe of market liquidity metrics to choose from.

<sup>19</sup>Although the magnitude of customer order flow is seldom directly observable, a number of important models, including the price-impact measure of Kyle [1985] and the *VNET* model of Engle and Lange [2001], require only a direction-of-trade indicator, which can be inferred (with error) from the sign of sequential price changes.

For example, turnover (trading volume divided by total outstanding) is a commonplace heuristic for liquidity in many markets, but its interpretation differs across markets. For bonds and equities, the denominator is simply the total amount issued; for exchange-traded futures, which lack a fixed issue amount, one might substitute open interest; but for over-the-counter markets, such as interest-rate swaps or foreign exchange, the choice of a plausible denominator and interpretation of the resulting measure is ambiguous.

#### *Market-invariant price impact*

Among the candidate metrics, we choose the invariant price-impact measure of Kyle and Obizhaeva [2014], which captures the change in market prices caused by a one-directional order flow (buy or sell) of a given size. This metric is both feasible and timely, and is also designed to support comparability and granularity. Equation 1 shows the reduced-form (and empirically calibrated) price-impact relationship.<sup>20</sup>

$$C(X) = \frac{\sigma}{0.02} \left[ \frac{8.21}{10^4} \left( \frac{W}{(0.02)(40)(10^6)} \right)^{-1/3} + \frac{2.50}{10^4} \left( \frac{W}{(0.02)(40)(10^6)} \right)^{1/3} \frac{X}{(0.01)V} \right] \quad (1)$$

In particular, equation (1) measures price impact as the market-specific (i.e., granular) expected daily volatility of returns,  $\sigma$ , normalized by a complicated scale factor (in square brackets) that adjusts for local price-level and expected volume conditions; the normalization yields a measure that should be invariant (i.e., comparable) across markets and over time. Here,  $C(X)$  is the trading cost as a response to a trading impulse of size  $X$ , where

---

<sup>20</sup>Kyle and Obizhaeva [2014, equations (70) and (71)] present two alternative versions of the price-impact measure, which differ in the functional form of the response of transaction costs to speculative order flow, which is allowed to be either linear, as in equation (1), or obey a square-root rule:

$$C(X) = \frac{\sigma}{0.02} \left[ \frac{2.08}{10^4} \left( \frac{W}{(0.02)(40)(10^6)} \right)^{-1/3} + \frac{12.08}{10^4} \left( \frac{X}{(0.01)V} \right)^{1/2} \right]$$

In our sample, the two versions produce qualitatively similar results, and we focus on the linear specification.



$V$  is expected daily trading volume (in shares or analogous units),  $X$  is a typical order size for the specific market, and  $W$  as the level of speculative activity (measured as price times expected volatility times expected volume). The first term inside square brackets is the portion of trading cost attributable to the bid-ask spread, and the second term (involving  $X$ ) is the price-impact component. To facilitate an intuitive interpretation of the final result,  $W$  is scaled by a factor  $W^* = (.02)(40)(10^6)$ , which is simply a benchmark  $W$  value for a hypothetical stock; similar scale factors are applied to the other terms in (1).

Significantly, the normalization factor in (1) embeds important structure, derived from theoretical first principles asserted to describe trading in speculative markets. The basic intuition of the invariance measure is that illiquidity reveals a market’s resilience, or lack thereof, to net speculative order flow. Speculative bets represent individual decisions to take on (or unload) risk; they tend to arrive at different rates in different markets, creating a phenomenon of market-specific business time defined by the pace of speculative trading.<sup>21</sup> Such bets reflect the market’s net risk-bearing capacity — long and short — which is the ultimate source of liquidity. Kyle and Obizhaeva [2014] argue that the bet arrivals can be approximated by a Poisson process with arrival rate  $\gamma$ , so that *expected* speculative order flow is proportional to calendar time (one unit of business time equals  $1/\gamma$ ). Similarly, the observed returns variance,  $\bar{\sigma}^2$ , can differ across markets for many reasons, but is assumed to be a constant multiple of an underlying, market-specific betting variance,  $\sigma^2$  (i.e., that is caused by speculative order flow as opposed to news-induced volatility, for example).

The invariance hypothesis is that, after normalizing by the local speculative capacity of the market — the amount of risk transferred per unit of business time — the arrival of a bet of dollar size  $PQ$  will generate a dollar price-impact distribution whose variance

---

<sup>21</sup>This is a recent addition to the literature on “time deformation” approaches to improving the empirical regularity of financial time-series, which stretches back at least to the pioneering work of Müller et al. [1990] on intradaily foreign exchange data. Other examples include Drost and Nijman [1993], and the ACD-GARCH model of Engle [2000]. The theoretical model of Easley and O’Hara [1992] envisions a subdivision of the trading day into equally spaced (in calendar time) intervals, but with an interval length that can vary across markets to accommodate local conditions.

depends only on the market-specific volatility conditions:  $PQ(\bar{\sigma}^2/\gamma)$ , or, equivalently in terms of standard deviation,  $PQ(\bar{\sigma}\gamma^{-1/2})$ . In other words, in equilibrium, overall speculative capacity allocates itself across the system to maintain this as an empirical relationship that is constant across markets (and over time).

Equation (1) is the final reduced-form result of: simple algebraic manipulations; the inclusion of a bid-ask spread component to expand from a pure price-impact relationship to a more general transactions cost function,  $C(X)$ ; and calibration of the remaining parameters against actual data. A factor of  $\gamma^{3/2}$  emerges naturally in these transformations as a product of: a linear scaling by business time from individual bets to observed daily volume; and the square-root scaling of the volatility of the price impact distribution noted above.<sup>22</sup>

In applying equation (1), there is no unambiguously right way to set the typical order size,  $X$ . An important consideration is to normalize  $X$  by trading activity in each market, to measure price-impact responses on a comparable scale across markets. A corollary requirement is to calibrate (1) to be consistent with the definition of  $X$ . The particular calibration in (1) assumes that order size is a constant fraction of average daily volume. This implies, for example, that the dollar size of the order should move in the same direction as dollar volume. A plausible alternate is to hold the dollar size of an order constant over time, so that the relative size of the order (as a fraction of volume) moves inversely with volume. As a robustness check, without re-estimating the parameters in (1), we recalibrated order size as a constant dollar value. The price-impact results were similar in magnitude, but noisier than for the calibration of  $X$  as a constant fraction of daily volume; the results presented below use the constant-fraction specification.<sup>23</sup> We follow Kyle and Obizhaeva

---

<sup>22</sup>This  $\gamma^{3/2}$  factor is the ultimate source of the curious exponents in equation (1). We can provide only a very basic sketch of the intuition underlying (1) in the limited space here. Kyle and Obizhaeva [2014] provide a detailed derivation of the algebraic manipulations that result in (1), along with numerous other insights and examples.

<sup>23</sup>Another possibility is to allow the size of the orders to adjust to market liquidity changes, in a manner more rigorously consistent with the equilibrium arguments in Kyle and Obizhaeva [2014]. For example,

[2014] in using the average trading volume over the preceding month (20 trading days) as a proxy for expected volume in (1). Similarly, we use the average realized volatility of daily returns over the preceding month as a proxy for expected volatility.<sup>24</sup>

### *Comparison to Other Approaches*

The invariance measure of Kyle and Obizhaeva [2014] is only one of several established market liquidity metrics that satisfies the four requirements set out above. To justify our choice, we compare the invariance measure (labeled INVL below) to a selection of other metrics that are acceptable under our criteria. In addition to satisfying the requirements, these alternatives were chosen to represent a diverse range of approaches:

- AMIH – This measure, defined by Amihud [2002] is based on the notion, originally advanced by [Amihud and Mendelson, 1986], that illiquidity should be priced and therefore should appear in returns. The basic equation is the daily absolute return,  $|R_{i,t}|$ , for security  $i$ , divided by daily volume,  $v_{i,t}$ :

$$\text{AMIH} = \frac{|R_{i,t}|}{v_{i,t}}$$

(Variations on this measure are also available, precalculated, for download from the NYU Volatility Institute (V-Lab) [2014].)

---

speculative order flow should be directly proportional to both overall liquidity and to the cube root of expected dollar volume, per Kyle and Obizhaeva [2014, equation (8)], so that order size should increase with liquidity, while decreasing as a fraction of daily volume. We are grateful to Pete Kyle for a helpful clarifying conversation around this issue. Because we are interested here in applying rather than testing or extending the Kyle and Obizhaeva [2014] model, we restrict attention to the calibrations in equations (1), together with a fixed dollar size for the order impulse,  $X$ .

<sup>24</sup>There are obviously more sophisticated ways to estimate conditional expected volatility, including the many GARCH specifications available in the literature; see Bollerslev [2010]. However, Kyle and Obizhaeva [2014, p. 31] note that using a more exacting ARIMA model in log volatilities produces quantitatively similar results. Again, because we are interested in applying rather than refining their model, we apply the simpler specification.

- LVOL – This is simply the logarithm of expected volatility,  $\hat{\sigma}_{i,t}$ :

$$\text{LVOL} = \ln(\hat{\sigma}_{i,t}),$$

where  $\hat{\sigma}_{i,t}$  is estimated as the standard deviation of daily returns over a rolling window ending at day  $t$ . We include this measure, because  $\hat{\sigma}_{i,t}$  plays such a prominent role as the leading term of equation 1.

- ROLL – Roll [1984] proposes to infer (approximately) the quoted spread from the assumption that the time-series of price changes is dominated by bid-ask bounce:

$$\text{ROLL} = 2\sqrt{-\text{cov}(\Delta p_{i,t-1}, \Delta p_{i,t})}$$

This is a workaround for the fact that pre-trade transparency is limited for many markets.

- KLAM – Kyle’s lambda, originally defined by Kyle [1985], is a commonly used price-impact measure. We calculate it as a cross-sectional average (across  $N$  firms) of estimated price-impact coefficients,  $\hat{\lambda}_i$ :

$$\text{KLAM} = \frac{\sum_{i=1}^N \hat{\lambda}_i}{N}$$

where the  $\hat{\lambda}_i$  values are calculated by regressing daily returns on signed dollar volume:

$$R_{i,t} = \hat{c}_i + \hat{\lambda}_i \cdot \text{Sgn}(t) \log(v_{i,t} p_{i,t}) + \epsilon_{i,t}$$

where  $\text{Sgn}(t) \in \{+1, -1\}$  indicates the direction of the trade (net customer buy or sell, respectively), and  $R_{i,t}$ ,  $v_{i,t}$ , and  $p_{i,t}$  are the return, trading volume and price of security  $i$  on day  $t$ .

We compare the four measures with the market invariant measure in equation 1, by estimating all five on the full universe of CRSP equities, Standard Industrial Classification 6 (SIC 6), between January 1986 and March 2014. We work with equities alone for this exercise, to reveal the behavior of the metrics over a longer time span. Large-scale systemic liquidity events like the crisis episode of 2008 are rare, and one of the purported advantages of the invariant approach is its comparability over time.<sup>25</sup> Understanding how the liquidity measures perform across diverse historical episodes is therefore an important exercise.

Table 1 presents simple linear correlations among the five series. Figure 5 plots the time series for all five measures over nearly 30 years, from 1986 to 2014. Basic sample statistics appear in Table 2. All five series respond strongly to the major liquidity events of 2008, and all the measures are positively correlated with each other, indicating that all the measures are tracking some facet of market liquidity. The invariance metric (INVL) is strongly correlated with each of the four other measures, but in no case is the correlation perfect. Interestingly, the strongest correlation is between Kyle’s lambda (KLAM) and log volatility (LVOL).

To facilitate visual comparison, the series are rescaled so that the in-sample maximum equals 1.00 and the minimum equals zero. The comparison offers support for the Kyle and Obizhaeva [2014] metric (INVL) as an approach to systemic liquidity monitoring. The most obvious pattern in Figure 5 is the pronounced secular trend in Kyle’s lambda (KLAM) and volatility (LVOL). The correlation is quite high between these measures, due to the regression equation that creates the  $\hat{\lambda}$  series, and so both appear to be picking up the “Great Moderation.” From a monitoring perspective, this instability in the measures

---

<sup>25</sup>Among the series in our sample, oil futures also have a long history, and the analysis of those series (as well as the other equities series) yields qualitatively similar results. The invariance measure registers a spike in illiquidity around the start of the first Gulf War in 1991, which does not show up significantly in equities markets. Kyle and Obizhaeva [2013] apply the invariant approach to five historical large liquidity events: the October 1929 stock market crash; the “Black Monday” crash of October 1987; a subsequent event in the futures markets, three days after the 1987 crash; the Société Générale rogue trader event in January, 2008; and the May 2010 “flash crash” in the futures markets.

Table 1: Correlations Among Liquidity Measures for Financial Equities

January 1986 – March 2014					
	AMIH	LVOL	ROLL	KLAM	INVL
AMIH	1.00	0.39	0.16	0.39	0.42
LVOL		1.00	0.23	0.92	0.75
ROLL			1.00	0.11	0.44
KLAM				1.00	0.66
INVL					1.00
Sources: CRSP, WRDS, OFR analysis					

is troubling, because it complicates comparisons across liquidity events that are separated in time, and makes interpretation of the signals context-dependent. The relative lack of skewness and kurtosis for these measures (see Table 2) is also an artifact of this trend. Due to the nonlinear nature of illiquidity, one expects the metrics to be skewed. On the other hand, a visual inspection of the series reveals that the Amihud and Roll measures are relatively noisy, with transient spikes that can be very large in scale. The other three metrics (Kyle’s lambda, volatility, and invariant price-impact) avoid this, likely because they all incorporate some form of moving-window estimation to smooth the series. In any case, the amplified noise-to-signal ratio reduces the usefulness of the Amihud and Roll measures as monitoring tools.

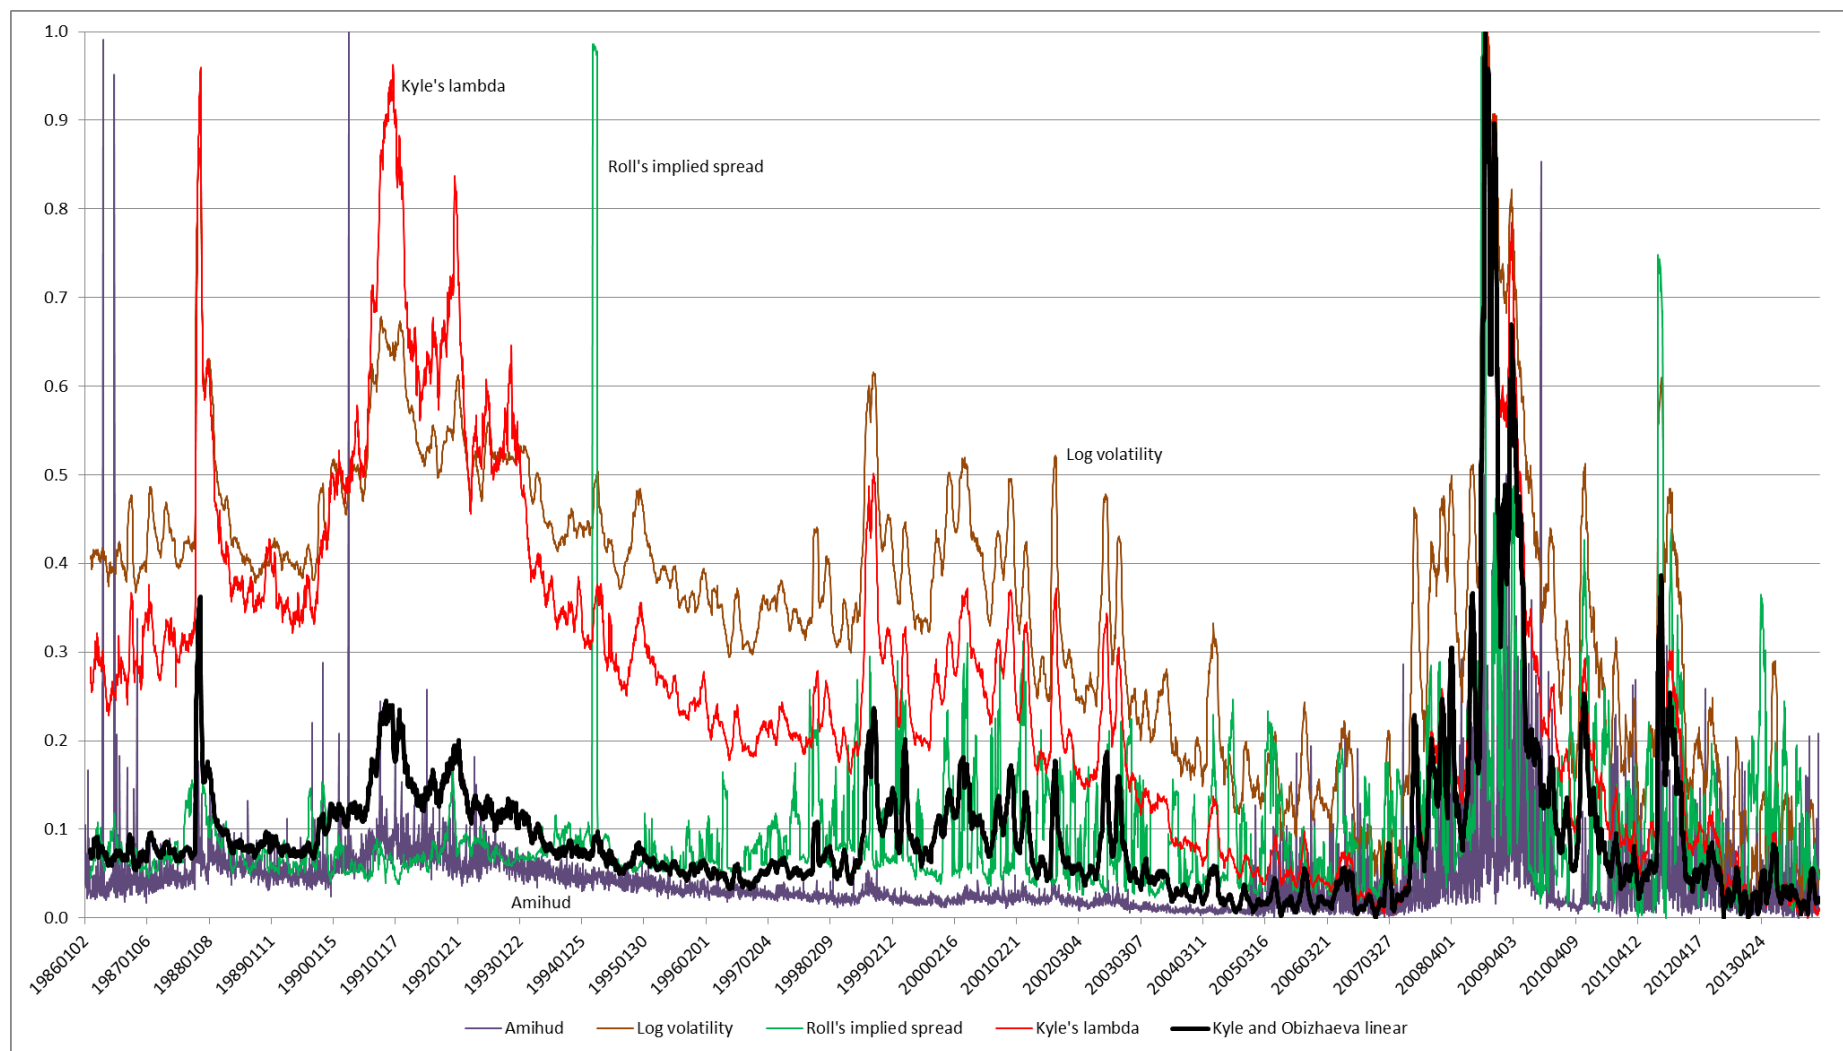


Figure 5: Five market liquidity measures for SIC 6 financial industry equities, January 1986 – March 2014,  
 Sources: CRSP, Bloomberg L.P., Mergent Inc., WRDS, St. Louis Federal Reserve Economic Data, OFR analysis

Table 2: Five Market Liquidity Measures for Financial Equities

January 1986 – March 2014					
	AMIH	LVOL	ROLL	KLAM	INVL
Mean	0.04330	0.35720	0.10243	0.26807	0.09706
Std. deviation	0.04803	0.16116	0.09747	0.19360	0.10050
Skewness	7.37033	0.26863	4.94018	1.09850	4.30274
Kurtosis	101.59150	0.52763	35.01501	1.23800	25.59540
Num. observations	7121	7121	7121	7121	7121
Sources: CRSP, WRDS, OFR analysis					

Overall, the comparison supports the selection of the invariant price-impact measure of Kyle and Obizhaeva [2014] for systemwide monitoring of market liquidity conditions. First, and most important, it satisfies the requirements of feasibility, timeliness, comparability, and granularity set out at the beginning of this section. In addition, when compared to several other commonly used market liquidity metrics, the invariance measure exhibits more consistency over time and is less subject to general noise and transient spikes.

## 2.2 Univariate Models of Latent Structure

The primary assumption underlying the hidden Markov chain (HMC) approach is that the liquidity in a specific financial market (as defined by a portfolio of publicly traded securities) jumps between distinct states (e.g., low to high liquidity), and then stays in that new state a some random period of time. The subsequent observed level of liquidity is a random deviation from the average liquidity level particular to the underlying current state. Although observed liquidity can occasionally fall between the average liquidity levels for two neighboring states, ultimately the persistence of the observed liquidity resolves the ambiguity to identify a single underlying state.

When liquidity from multiple financial markets is considered, we augment the liquidity models for each individual market with an add-on hierarchical model that can explain,



in part, periods of coordination in which a large subset of the financial markets exhibit similar liquidity patterns. The hierarchical portion of the model, which links individual liquidity models together, is an add on in the sense that it does not feed back into the individual level liquidity models. Instead it allows us to determine whether summary series describing the broader financial markets and economy are related to liquidity states across multiple markets. Uncovering such a relationship offers a framework for understanding and potentially predicting systemwide liquidity stress (by either lagging the summary series and/or predicting the underlying dynamics of the summary series).

Among the many available modeling approaches, we choose to focus on regime shifting models, as they naturally account for sudden jumps in liquidity, which is an empirical characteristic of financial crises and they allow us to remove slow moving, diffusion type dynamics from our multivariate model, which aids in our ability to predict. We briefly explore more traditional models, such as a vector autoregressive (VAR) model, but find that the multi-collinearity in the multivariate data (without doing some type of dimension reduction) results in insights from the VAR models that are of little value.<sup>26</sup> Section 5 provides some summaries of the performance of a VAR model. In an acknowledgment that we observe slow moving diffusion dynamics in liquidity measurements, in conjunction with sudden jumps, we introduce a version of the univariate model that allows for autoregression around a level, which switches between different levels of attraction.<sup>27</sup>

The analysis of latent structure begins with a univariate, hidden Markov chain for each financial market, where liquidity is a random deviation from a latent value associated with each state of the hidden Markov chain. We consider two variations of random deviations: independent deviations around an average level; and deviations around a value that mean-

---

<sup>26</sup>See Koop and Korobilis [2010] for a discussion of Bayesian approach to VAR models.

<sup>27</sup>An alternative ad hoc approach to the regime-switching model would be to define regimes by breaking the data into percentile regions — e.g., lower quartile, interior quartiles, and top quartile. However, such an approach would not provide the persistence of staying in a state, allowing instead for sudden spurious jumps. More importantly, it would not find the long-term levels of attraction that exist naturally in the data.

reverts around an average level. Initially, we assume the dynamics of these models are unrelated; we later propose a hierarchical (multiple-market) model where systemwide market summaries explain the states identified by the collection of univariate hidden Markov models.

Liquidity measurements over  $T$  periods,  $y_i = (y_{i1}, \dots, y_{iT})^T$ , for market  $i$  are assumed to be random normal deviations around a dynamic, latent level of liquidity  $\theta_i = (\theta_{i1}, \dots, \theta_{iT})$ , or

$$y_i = \theta_i + \epsilon_i,$$

where  $\epsilon_i \sim N(0, \sigma_i^2 I_T)$ , and  $I_T$  is a  $T$ -dimensional identity matrix. For the first version of the model (the HMC-only version), the latent level  $\theta_i$  is one of  $K$  levels, each of which represents a different level of liquidity or state for each market specific, discrete-time hidden Markov chain  $D_i$ , or

$$\theta_i = F_i \bar{\theta}_i,$$

where  $\bar{\theta}_i$  is a  $K \times 1$  vector. Each element represents the average level of liquidity for the  $k^{th}$  state of  $D_i$ , or

$$F_i(t, k) = I\{D_{it} = k\}$$

and  $I\{\}$  is an indicator function equaling either 0 or 1.

The HMC version is typically sufficient to identify structural shifts in liquidity patterns. However, there are some markets for which the local variation in the level of liquidity supports an overly large number of hidden states. In these cases, we use a mean-reverting version of the hidden Markov chain model. For this second version of the model (the

mean-reverting hidden Markov chain, MRHMC), the latent level  $\theta_i$  mean-reverts around the average level associated with the state of  $D_i$ , or for  $t = 2, \dots, T$ ,

$$\Delta\theta_i = \gamma_i((\theta_i)_{-T} - (F_i\bar{\theta}_i)_{-1}) + (\xi_i)_{-1}, \quad (2)$$

where  $()_{-1}$  indicates that the first element and  $()_{-T}$  indicates that the last element of the vector  $()$  have been removed, and

$$\Delta\theta_{it} = \theta_{it} - \theta_{it-1}.$$

For  $t = 1$  let

$$\gamma_i\theta_{i1} = \gamma_i\bar{\theta}_{i1} + \xi_{i1},$$

where  $\xi_i \sim N(0, w_i I_T)$ .

We require  $0 < \gamma_i \leq 1$ , which ensures that  $\theta_i$  is stationary and increases the variance of  $\theta_{i1}$ , allowing the starting value of  $\theta_i$  to be relatively vague. Alternatively, (2) can be rewritten as,

$$L_i\theta_i = \gamma_i F_i \bar{\theta}_i + \xi_i, \quad (3)$$

where  $L_i$  is a sparse  $T \times T$  matrix with zeros except for the following elements,  $L_i(j, j) = 1$  and  $L_i(j, j-1) = \gamma_i - 1$  for  $j > 1$  and  $L_i(1, 1) = \gamma_i$ .

For both versions of the model, the dynamics of  $D_i$  are given by an initial probability density  $\nu_i$ , a  $K \times 1$  vector, and a transition probability density  $P_i$ , a  $K \times K$  matrix. Given

a realization of  $D_i$ , its density is given by

$$f(D_i) = \nu(D_{i0}) \prod_{t=1}^T P_i(D_{it-1}, D_{it}).$$

We assume conjugate priors for  $\sigma_i^2$  and  $w_i$  (inverted Gamma),  $\nu_i$  and each row of  $P_i$  (Dirichlet) and  $\bar{\theta}_i$  and  $\gamma_i$  (truncated normal). In addition, we use subjective priors based on initial conditional maximum likelihood estimates of summaries of the data, to ensure that the filtered HMC model can clearly distinguish between the dynamics of the hidden Markov chain and the dynamics of the latent value  $\theta_i$ .

#### *Full Conditional Distributions HMC Model*

We use Markov chain Monte Carlo (MCMC) analysis to infer parameter values for both of these univariate models and the multivariate model built on these univariate models.<sup>28</sup> The full conditional densities used in the MCMC analysis for the HMC model are:

$$\bar{\theta}_i | - \sim N \left( \left( \frac{1}{\sigma_i^2} F_i^T F_i + \frac{1}{\tau_{\bar{\theta}_i}^2} I_K \right)^{-1} \left( \frac{1}{\sigma_i^2} F_i^T y_i + \frac{\mu_{\bar{\theta}_i}}{\tau_{\bar{\theta}_i}^2} \bar{\theta}_i \right), \left( \frac{1}{\sigma_i^2} F_i^T F_i + \frac{1}{\tau_{\bar{\theta}_i}^2} I_K \right)^{-1} \right) I\{\bar{\theta}_{i1} < \dots < \bar{\theta}_{iK}\};$$

where  $-$  represents everything else remaining in the model, and

$$\frac{1}{\sigma_i^2} | - \sim \text{Gamma} \left( \text{shape}_{\sigma_i^2} + \frac{T}{2}, \text{scale}_{\sigma_i^2} + \frac{1}{2} (y_i - F_i \bar{\theta}_i)^T (y_i - F_i \bar{\theta}_i) \right).$$

Realizations of the hidden Markov chain  $D_i$ , conditional on the remaining parameters and data, are generated following the filter-forward, sample-backward approach commonly used with discrete-time Hidden Markov chains.<sup>29</sup> For completeness, the filter-forward

---

<sup>28</sup>For a description of MCMC methods, see Brooks et al. [2011] and Gelman et al. [2013].

<sup>29</sup>See Baum et al. [1970]. Cappé et al. [2005] describe these in a more general continuous-time framework.

equations for the HMC model are given below,

$$f(y_{it}|-,\mathcal{F}_{it-1}) = \sum_{k=1}^K f(y_{it}|D_{it}=k,-,\mathcal{F}_{it-1}) f(D_{it}=k|-,\mathcal{F}_{it-1}), \quad (4)$$

where  $\mathcal{F}_{it} = \{Y_{i1}, \dots, Y_{it}\}$ ; and by

$$f(D_{it}=k|-,\mathcal{F}_{it}) = \frac{f(y_{it}|D_{it}=k,-,\mathcal{F}_{it-1}) f(D_{it}=k|-,\mathcal{F}_{it-1})}{f(y_{it}|-,\mathcal{F}_{it-1})}. \quad (5)$$

Specifying a vague initial state probability, e.g.,

$$f(D_{i0}=k|-,\mathcal{F}_{i0}) = \frac{1}{K},$$

completes the forward recursion. The key equation for the backward sampling is given by the density of the hidden Markov chain, conditional on all of the data, or

$$f(D_{iT-t}=k|-,\mathcal{F}_{iT}) = \sum_{j=1}^K \frac{f(D_{iT-t+1}=j|D_{iT-t}=k,-,\mathcal{F}_{iT-t}) f(D_{iT-t}=k|-,\mathcal{F}_{iT-t})}{f(D_{iT-t+1}=j|-,\mathcal{F}_{iT-t})} f(D_{iT-t+1}=j|-,\mathcal{F}_{iT-t+1}). \quad (6)$$

Given these formulas, generating a realization is straightforward: i) calculate the forward filter; ii) generate a sample for  $D_{iT}$  from (5), with  $t = T$ ; and iii) recursively calculate  $f(D_{iT-t}=k|-,\mathcal{F}_{iT})$ , conditional on all of the draws  $(D_{iT}, \dots, D_{iT-t+1})$  using (6) and use this to generate a sample for  $D_{iT-t}$ . Given a realization of  $D_i$ , the full conditional distribution for each row of the transition probability is given by

$$P_i(j, :)|- \sim \text{Dirichlet}(\alpha_{i1} + n_{ij1}, \dots, \alpha_{iK} + n_{ijK}),$$

where  $\alpha_{ijk}$  is the prior associated with  $D_i$  jumping from state  $j$  to  $k$  and  $n_{ijk}$  is the actual

number of times that the current realization of  $D_i$  jumps from state  $j$  to state  $k$ . A similar full conditional density exists for  $\nu_i$ , but this is inconsequential, because the backward recursion dominates the *a priori* initial state. It is important to note that the MRHMC model is disentangling two dynamics, the dynamics of  $D_i$  and  $\theta_i$ . In practice, we found that the model required a strong priors on the dynamics of  $D_i$  to obtain meaningful distinctions between these two dynamics. Setting  $\alpha_{ikk}$  to a sufficiently large value, suggesting *a priori* that the hidden chain is persistent, results in a clean separation of these two competing dynamics.

#### *Full Conditional Distributions MRHMC Model*

There are similarities between some of the full conditional densities of the MRHMC model and the HMC model. The full conditional density for  $P_i$  is unchanged, while the full conditional density for  $\bar{\theta}_i$  is obtained by replacing  $y_i$  with  $\frac{1}{\gamma_i}L_i\theta_i$  and  $\sigma_i^2$  with  $\frac{w_i}{\gamma_i^2}$ . The full conditional density for  $\frac{1}{\sigma_i^2}$  is obtained by replacing  $F_i\bar{\theta}_i$  with  $\theta_i$  and the full conditional density for  $D_i$  is obtained by replacing the likelihood  $f(y_{it}|D_{it} = k, -, \mathcal{F}_{it-1})$  used in (4) and (5) with  $f(\theta_{it}|D_{it} = k, -, \theta_{it-1}, \dots, \theta_{i1})$ . The remaining full conditional densities for the MRHMC model are as follows:

$$\frac{1}{w_i}|- \sim \text{Gamma}\left(\text{shape}_{w_i} + \frac{T}{2}, \text{scale}_{w_i} + \frac{1}{2}(L_i\theta_i - \gamma_i F_i\bar{\theta}_i)^T(L_i\theta_i - \gamma_i F_i\bar{\theta}_i)\right);$$

$$\gamma_i|- \sim N\left(\Sigma_i\left(\frac{1}{w_i}A_i^T\Delta\theta_i + \frac{\mu_{\gamma_i}}{\tau_{\gamma_i}^2}\right), \Sigma_i\right)I\{0 < \gamma_i \leq 1\};$$

where

$$\Sigma_i = \left(\frac{1}{w_i}(A_i^T A_i + (\theta_{i1} - \bar{\theta}_{iD_{i1}})^2) + \frac{1}{\tau_{\gamma_i}^2}\right)^{-1} \text{ and } A_i = ((\theta_i)_{-T} - (F_i\bar{\theta}_i)_{-1}),$$

and

$$\theta_i | - \sim N \left( \left( \frac{1}{w_i} B_i + \frac{1}{\sigma_i^2} I_K \right)^{-1} \left( \frac{\gamma_i}{w_i} B_i (L_i^{-1} F_i \bar{\theta}_i) + \frac{1}{\sigma_i^2} y_i \right), \left( \frac{1}{w_i} B_i + \frac{1}{\sigma_i^2} I_K \right)^{-1} \right),$$

where  $B_i$  is a  $T \times T$  matrix given by

$$B_i = (L_i^{-1} (L_i^{-1})^T)^{-1}.$$

An alternate approach for sampling  $\theta_i$  and  $\gamma_i$ , conditional on  $D_i$ , is to treat them as a discrete-time, dynamic linear model and use a filter-forward, sample-backward strategy like the Kalman filter.<sup>30</sup> Although we explored a filter-forward, sample-backward approach, we found that this was not as stable as the regression-based approach detailed above. Obviously, one disadvantage of the regression approach is the need to calculate  $B_i$ , which requires the inversion of a  $T \times T$  matrix, which can become computationally prohibitive as  $T$  becomes large. Fortunately, the form of  $L_i$  results in a banded matrix for  $B_i$ , where all elements are zeros except for the main diagonal and neighboring diagonals. In addition, the non-zero elements are functions of  $\gamma_i$ ; to be explicit,

$$B_i(j, j) = \begin{cases} 1, & \text{if } j = T \\ \frac{1}{2} + 2 \left( \frac{1}{2} - \gamma_i \right)^2, & \text{otherwise} \end{cases}$$

and

$$B_i(j, j-1) = B_i(j-1, j) = \gamma_i - 1.$$

---

<sup>30</sup>See Kálmán [1960] for the original reference, and Carter and Kohn [1994] and De Jong and Shephard [1995] for MCMC-based inference approaches.

## 2.3 Hierarchical Model

The hierarchical add on model runs a collection of HMC or MRHMC models in parallel, one for each of the  $N$  markets under consideration. Each sweep of the MCMC algorithm generates a realization of the latent hidden Markov chain for each market, resulting in a collection of realizations  $(D_1), \dots, (D_N)$ . For each realization, every individual time point can be viewed as a draw from a multinomial distribution that is driven by a set of time-varying covariates  $x_t$ —the summary series.<sup>31</sup> Conditional on a current realization of the hidden Markov chain, the add on portion of the model is a multinomial probit model, where

$$f(D_{it} = k) = f(\tilde{z}_{itk} > \tilde{z}_{itl}, l \neq k), \quad (7)$$

and  $\tilde{z}_{it}$  is multivariate normal or

$$\tilde{z}_{it} \sim N(\tilde{\beta}x_t, \tilde{\Sigma}).$$

We follow the approach of McCulloch and Rossi [1994], which builds on Albert and Chib [1993], for dealing with the identification issues that arise in using a Bayesian approach for the multinomial probit model. A related, alternative approach is discussed in McCulloch et al. [2000]. The additive identification is overcome by forcing the latent value for state 1 to always be zero. This is done by defining  $z_{it}$  as follows,

$$z_{itk} = \tilde{z}_{itk} - \tilde{z}_{it1},$$

---

<sup>31</sup>Possible alternatives to the multinomial probit include an ordered probit model. The main difference would involve the interpretation of parameters. Although the ordered probit parameters would reveal the overall impact of a summary variable — for example, whether higher levels result in more or less liquidity — they would not readily reveal the levels of the summary series that lead to moderate levels of liquidity. We chose the multinomial probit model, because it would be more difficult under an ordered probit to distinguish systemwide conditions associated with high, medium and low levels of liquidity, while differences in estimation and empirical performance between the two approaches should be minimal.



which results in (7) becoming

$$f(D_{it} = k) = \begin{cases} f(0 > z_{itl}l \neq k), & \text{if } k = 1 \\ f(z_{itk} > \max(0, z_{itl}), l \neq k, l > 1), & \text{if } k > 1 \end{cases}$$

where  $z_{it} \sim N(\beta x_t, \Sigma)$ , and  $\beta$  is a  $(K - 1) \times p$  matrix, where  $p$  is the number of summary series, including an intercept. The scale identification is overcome by restricting

$$\Sigma_{1,1} = 1.$$

We assume conjugate priors for  $\beta$  and  $\Sigma$ . Following McCulloch and Rossi [1994], we sample  $\beta$  and  $\Sigma$  from the unconstrained full conditional densities using Gibb samplers and then rescale by dividing these draws by  $\Sigma_{1,1}$ , which enforces the above constraint.

The hierarchical portion of the model is considered to be an “add on” to the model because the distribution of the hidden Markov chains  $(D)_1, \dots, (D)_N$  does not depend on the multinomial probit probabilities or, more to the point, they do not depend on the summary series. For the distribution of the hidden Markov chains to depend on the summary series, we would need to model the transition between the latent liquidity states (as opposed to modeling the states themselves) as multinomial random variables conditional on the summary series; we leave this task for future research. Instead, the “add on” model summarizes the relationship between latent states and the summary series, acting as a supplemental analysis that describes how the latent liquidity states related to the summary series, but makes no assumption about nor gives any insights into how the summary series affect the dynamics of the latent states.<sup>32</sup> The basis of the relationship between the hidden

---

<sup>32</sup>A less sophisticated approach would be to save a realization of each hidden Markov chain from the MCMC analysis, and then calibrate a multinomial probit model for this collection of realizations. Repeat this multiple times, each time with a different set of realizations obtained by stopping the MCMC analysis at a random time, which would result in a set of multinomial probit parameter estimates, one for each set of realizations, and then average the parameters estimates from all of these analysis. Our approach is

states and the summary series is determined by the portions of the hidden Markov chains which are relatively stable (i.e., have a high probability of being in one of the states), which holds for large portions of time over each of the markets that we are considering. Portions of the hidden Markov chain that tend to switch states (have a probability that is distributed between two or more states) have less impact as the hidden Markov chains alternate between these competing states during the analysis and require estimates of  $\beta$ , which can reasonably accommodate this oscillation.

The fact that the hidden Markov chain can switch states during the analysis presents a technical challenge. When the hidden chain changes state, the latent values from the multinomial probit model, i.e., the  $z_{it}$ , must change to match their likelihood function. For example, assume that chain  $i$  at time  $t$  changes from state 2 to state 3; then  $z_{it3}$  must become positive and  $z_{it2}$  must be less than  $z_{it3}$ ). In practice, we found that when a hidden Markov chain changes state, we can sample from the truncated, full conditional density of each latent variables to impose the new ordering, but doing this once is typically not sufficient to provide a stability for the estimate of  $\beta$  and  $\Sigma$ . These stability issues can be overcome by drawing a small number of samples (on the order of a few dozen) of all of the related latent variables (e.g., draw repeatedly from  $z_{it}$  when a new ordering constraint is imposed by the change in state).

### 3 Data

We measure market liquidity on a daily basis across 33 markets, covering thousands of individual securities in four different asset classes. One important goal of casting a wide net across a diverse sample is to improve the chances of identifying emerging risks in liquidity, since it is difficult to assert a priori which market sector(s) will be affected first in an episode

---

more elegant as it updates the parameters of the multinomial probit model with each sweep of the MCMC analysis.

of illiquidity. Similarly, a broad panel should help in discerning significant patterns among the markets being monitored, as we map between local markets and system-level conditions. Finally, we hope that our mixing of several distinct asset classes in this analysis serves as an example of how to further expand the scope of the sample in subsequent research.

Specifically, our initial dataset includes the following instruments:

- All U.S. equities, January 1986 – March 2014, from CRSP, which provides comprehensive coverage of security price, return, and volume data for the NYSE, AMEX, and NASDAQ stock markets. CRSP also provides the Standard Industry Classification (SIC) for each security.
- All U.S. corporate bonds, July 2002 – March 2014, from TRACE, the Financial Industry Regulatory Authority’s (FINRA) real-time price dissemination service for the over-the-counter bond market. It provides transaction data for all eligible corporate bonds, which include investment grade and high-yield debt; we use the public TRACE database in this analysis.<sup>33</sup>
- West Texas Intermediate (WTI) light sweet crude oil futures, January 1986 – March 2014, from the New York Mercantile Exchange, the world’s largest-volume futures contract traded on a physical commodity. We collected data for contracts with expirations from one-month to six-months from Bloomberg.
- S&P 500 market volatility index (VIX<sup>®</sup>) futures, April 2004 – March 2014, from the Chicago Board Options Exchange. This is a pure play contract on implied volatility designed to reflect investors’ view of future (30-day) expected stock market volatility;

---

<sup>33</sup>We apply the heuristics of Dick-Nielsen [2009] to scrub the TRACE data. There is a separate “enhanced” version of the TRACE database, FINRA [2009], which does not truncate large trades, but which FINRA publishes only with a lag. We map TRACE bond identifiers (6-digit CUSIP codes) to the issuing firm’s SIC code, derived from CRSP, Mergent, or Bloomberg; approximately 2 percent of the bonds in sample could not be mapped, and were dropped from the analysis.

we collected data for contracts with expirations from one-month to nine-months from Bloomberg.

For the hierarchical analysis, we augment this with the systemwide market summaries detailed in Table 3 below.

Our primary analysis of the liquidity measures starts in 2004, when all series are available. We also provide some secondary comparisons of the longer-term performance of price-impact measures for equities and WTI futures, extending back to 1986. We grouped both the CRSP equities and TRACE corporate bonds data into portfolios based on one-digit SIC codes. For both bonds and equities the SIC portfolios cover SIC codes 0 through 8.<sup>34</sup> This clustering into portfolios reduces the dimensionality of the analysis and presentation of results. In the case of corporate bonds, the combination into portfolios is a practical necessity for the calculation of returns and volatility, because the trading of individual issues in this market is far too thin.<sup>35</sup>

We track the VIX<sup>®</sup> and WTI futures at the level of their relative maturity date, starting with the front-month contract. Actual calendar maturities follow a sawtooth pattern, as expiry dates gradually approach and abruptly transition to the next contract as expiration occurs. For both VIX<sup>®</sup> and WTI, and for the futures market generally, the near-dated contracts are usually more actively traded than the longer-maturity futures. There is no official longest maturity, but many possible long-dated contracts simply never trade. For the VIX<sup>®</sup> futures, we draw the line at nine different securities from the front month out to nine months forward. For the WTI futures, we use six different securities from the front month out to six months forward.

---

<sup>34</sup>The miscellaneous category (SIC 9, government establishments) is very lightly populated for both TRACE and CRSP, and did not provide sufficient observations for reliable analysis.

<sup>35</sup>Even with grouping into portfolios, there are numerous missing values in the time series of corporate bond activity. For calculating returns and volatility, we require that the most recent lagged observation be no older than a week (5 trading days). We experimented with higher and lower thresholds, out to 20 trading days, without a significant qualitative impact on the results.

## 4 Liquidity Regimes

In our initial analysis we estimated each price impact series independently, using both the hidden Markov chain (HMC) and the mean reverting hidden Markov chain (MRHMC) models. Although there is no coordination between the dynamics of the latent liquidity states across markets for this initial analysis, we find surprising consistency in the dynamics of market liquidity across all of these markets. Despite these common features, we also find interesting differences across the various markets in the lead-up to the recent crisis and in its aftermath. We formally explore these difference using the hierarchical model, which allows us to link the latent liquidity states (from multiple markets) together with a collection of summary series. This provides a framework for assessing the contribution of the summary series.

### 4.1 Individual Market Liquidity

We start by considering the performance of the two competing univariate models (HMC and MRHMC) and provide evidence that there are essentially three different liquidity regimes across these different markets. Then we report aggregate summaries based on these models.

#### *Performance of Models*

Both the HMC and MRHMC models identify interesting liquidity regimes within the Kyle-Obizhaeva price-impact data over the various markets that we considered, but that their relative performance depends on the amount of local variability of the liquidity in each state. The simpler HMC model can readily identify the three liquidity regimes for all of the equity markets; see Figure 6. In some cases, such as the front month of the WTI contract, the MRHMC model performed slightly better than the HMC model using

standard Bayesian model-choice tools; see Figure 7.

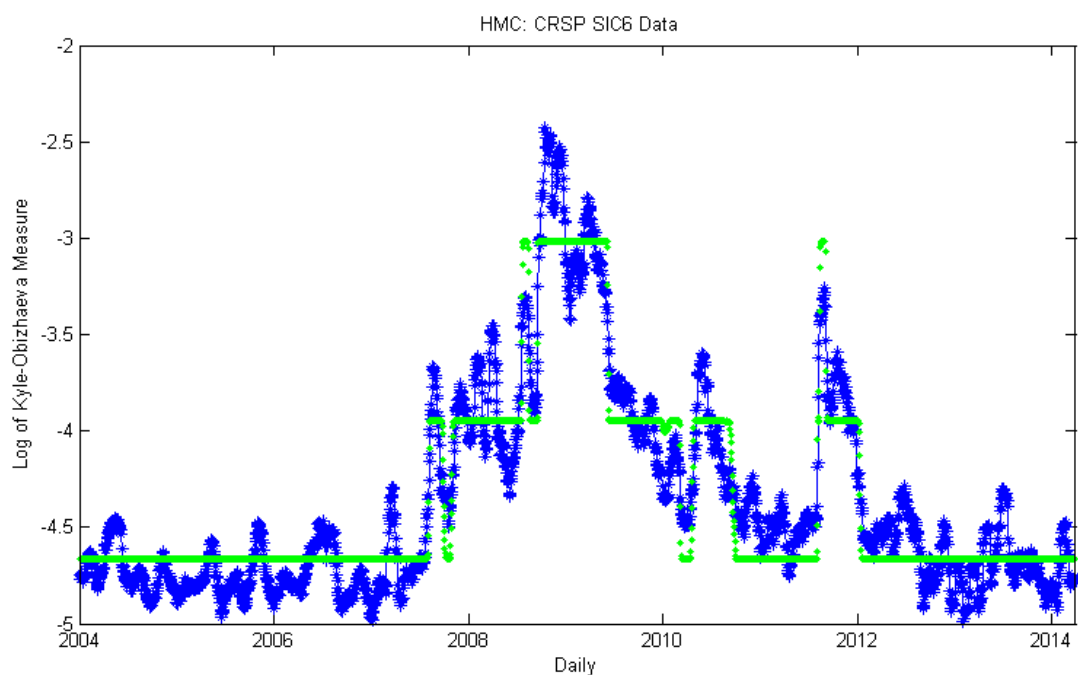
As a tool for policy makers, it is important to have a model that is parsimonious with respect to the number of states.<sup>36</sup> Therefore, we adopted a prior on the model space based on the number of hidden states, giving a high penalty for increased complexity; this resulted in three states being preferred for most markets. Alternative priors, with a smaller penalty for complexity would support a high number of latent states (typically in excess of 10 to 15); a visual inspection indicated that increasing the number of states essentially fractured the mid-liquidity state into a larger number of substates.

To give an initial graphical summary of the liquidity dynamics identified by the HMC model across 33 markets, we labeled these states the: (i) low, (ii) intermediate, and (iii) high price-impact states for each series, where high price impact means low liquidity, and vice versa. This analysis resulted in a daily estimate of the probability that each market was in each of these three unobserved states. Figure 8 presents the cross-sectional averages across the 33 series of these three probabilities, which must add up to one. Red indicates the likelihood of high price impact, and blue indicates low price impact; yellow is the intermediate state.

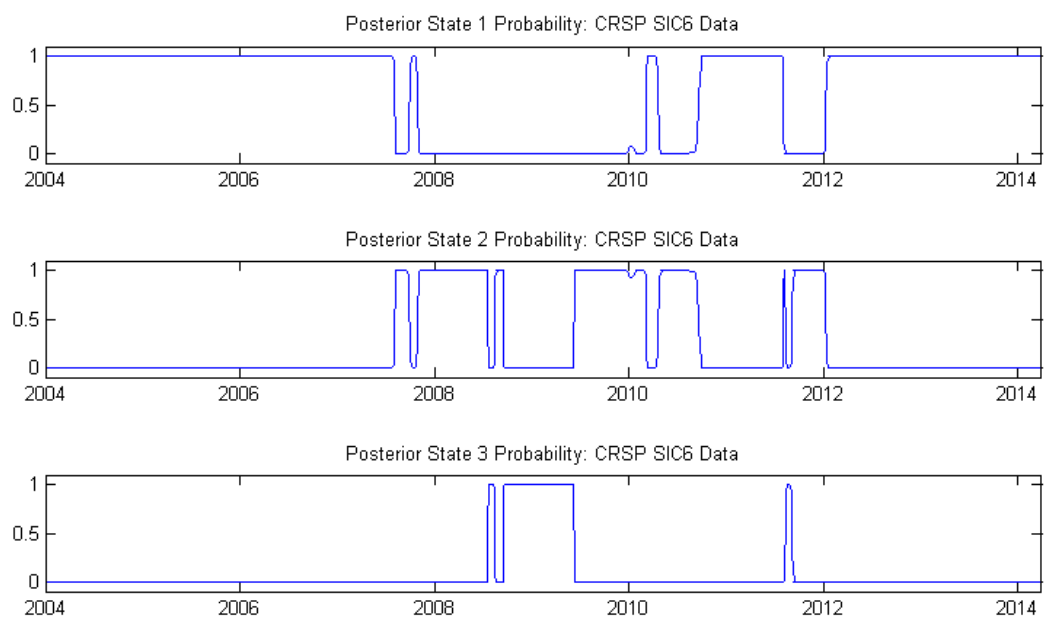
While there was diversity in market liquidity for these 33 series, there were also periods of common behavior. For example, in August 2011, the downgrade of U.S. Treasury debt by Standard & Poor's coincided with ongoing fiscal weakness in several eurozone countries and the initiation of the Occupy Wall Street movement to produce a sharp, but ultimately transient, spike in the probability of the low-liquidity (high price-impact) state. Similarly, the liquidity crisis after the failure of Lehman Brothers is plainly visible as the deep and more persistent spike in September 2008, preceded by a series of pronounced foreshocks

---

<sup>36</sup>Policy makers are the primary audience for this model, although we anticipate market participants will find it valuable as well. Because policy makers have to make decisions based on extreme market conditions, we felt that selecting three states was optimal, with state three being a crisis state. Of course it is straightforward to redo the entire analysis if policy makers or market participants feel that more or fewer states would provide a more useful insights.

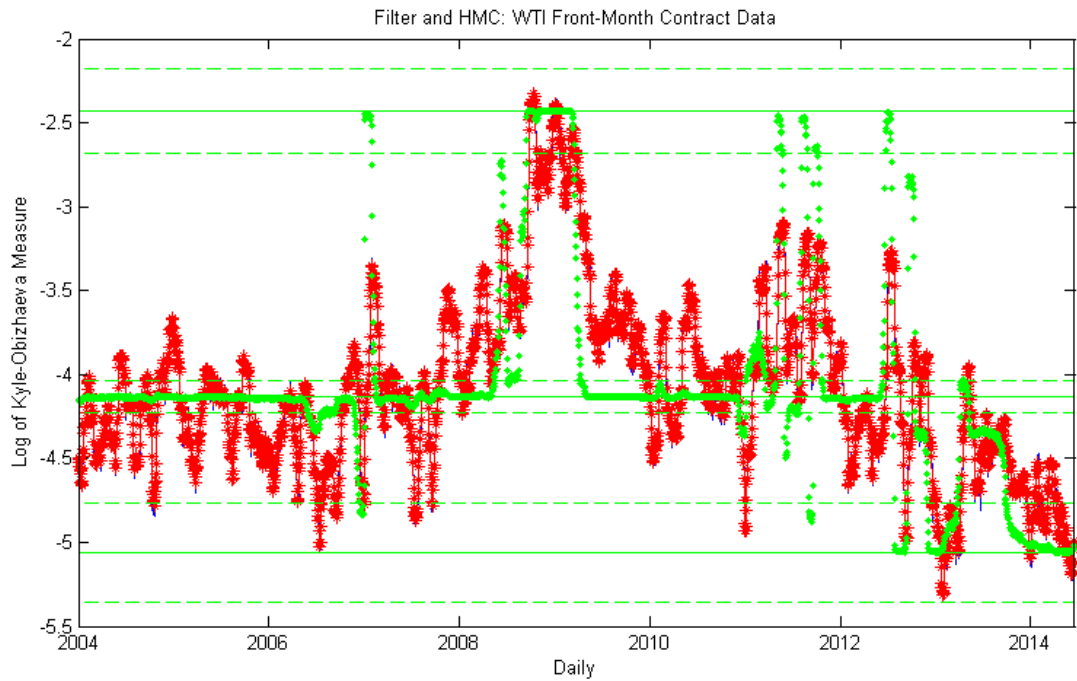


(a) Price-impact measure & level of attraction

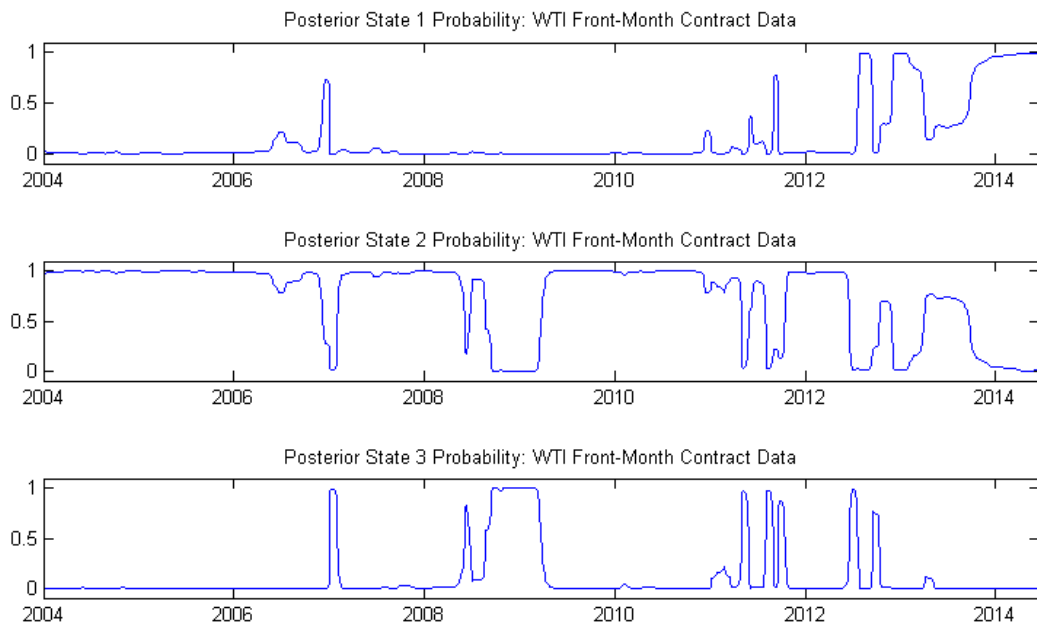


(b) HMC state probabilities

Figure 6: Equities, SIC 6, Kyle-Obizhaeva Measure and HMC Estimates  
Sources: CRSP, WRDS, OFR analysis



(a) Price-impact measure, mean-reversion level, & level of attraction



(b) HMC state probabilities

Figure 7: WTI futures, front month contract, Kyle-Obizhaeva Measure and MRHMC Estimates

Sources: Bloomberg L.P., OFR analysis



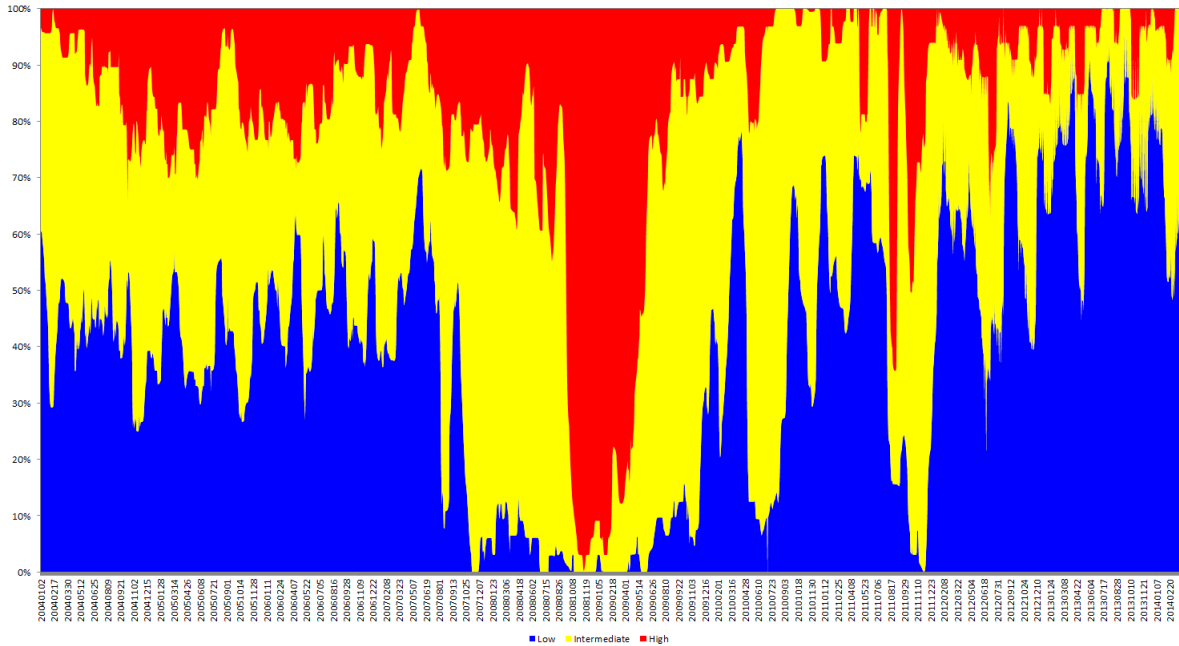


Figure 8: Probabilities of price-impact states, averaged across 33 markets, April 2004 – March 2014

Sources: CRSP, Bloomberg L.P., Mergent Inc., WRDS, FINRA, OFR analysis

over the course of the year.

Figure 9 condenses the trivariate time series of Figure 8 into univariate daily color codes by mixing the three primary colors of Figure 8 as a linear combination of RGB color vectors, weighted by their respective state probabilities on each day (black indicates a missing value). On any given day, one state, and therefore one color, tends to dominate the sample. Figure 9 shows a stacked sampling of such ribbons of daily data for four representative series. Figure 10 shows a similar stacked sample, group by asset class, for

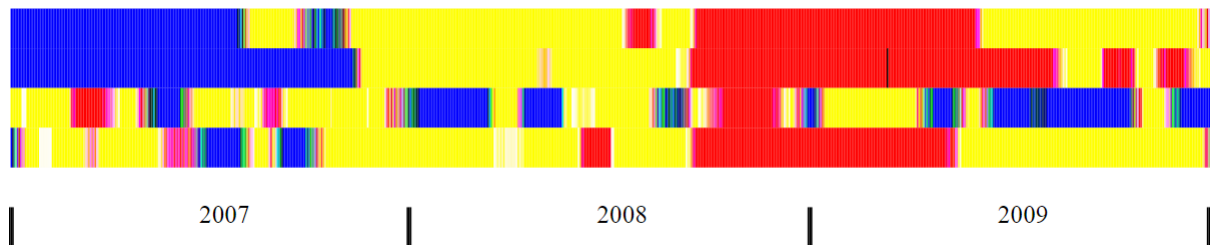


Figure 9: Daily Price-Impact Probabilities for Select Markets, Top to bottom: Equities SIC 6, Bonds SIC6, WTI futures, VIX<sup>®</sup> futures

Sources: CRSP, Bloomberg L.P., Mergent Inc., WRDS, FINRA, OFR analysis

all 33 series, covering the full sample period. The ribbon charts illustrate that the equity markets and the VIX<sup>®</sup> index responded strongly and immediately to the funding market distress in August 2007, but WTI futures did not. Throughout the 2007-09 crisis window, VIX<sup>®</sup> liquidity was more persistently stressed compared to the other asset classes in the sample. Corporate bond liquidity for series SIC 4 and SIC 8 took longer to recover from the elevated illiquidity levels of the crisis episode. Consistent with the increased uncertainty about the stocks in the financial sector, throughout late 2007 and 2008, remained depressed as we would expect, prior to the crisis liquidity in financial stocks remained depressed. Two other key insights from this analysis are that the liquidity implications of the Lehman Brothers failure were felt broadly for an extended period and that hints of illiquidity foreshocks existed in some markets, including financial stocks (SIC 6) and certain bond sectors, that may ultimately help in crafting liquidity forecasts.

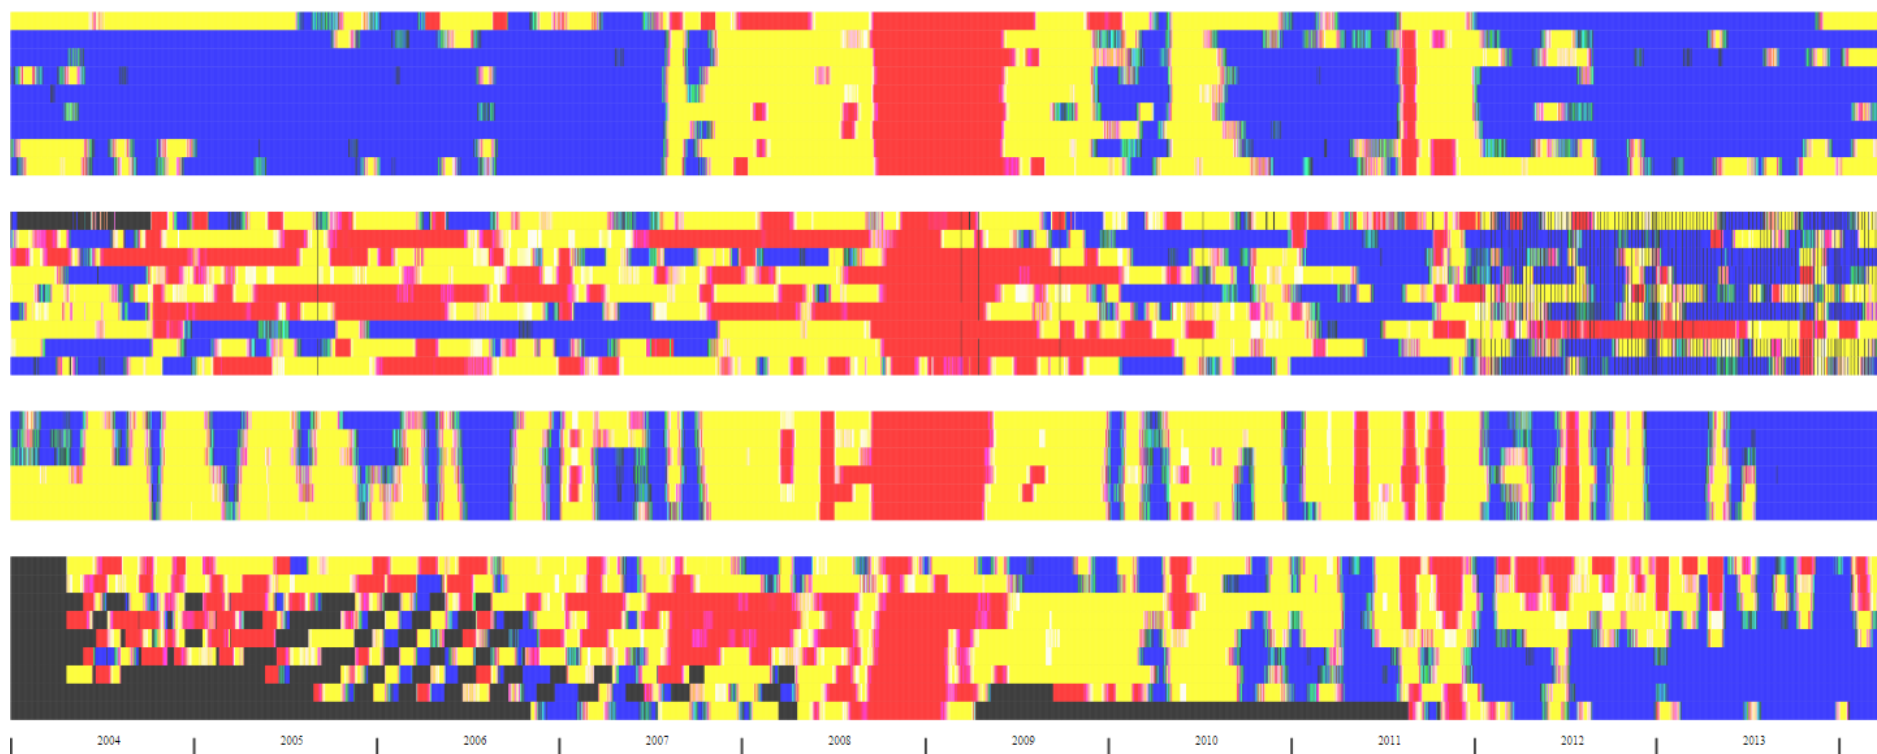


Figure 10: Daily Price-Impact Probabilities across all 33 Markets,  
 Top to bottom: Equities SIC 0-8, Bonds SIC 0-8, WTI futures, VIX<sup>®</sup> futures  
 Sources: CRSP, Bloomberg L.P., Mergent Inc., WRDS, FINRA, OFR analysis

## 4.2 Explaining Liquidity Regimes

There appear to be strong relationships between changes in the level of liquidity and a number of summary series. Although it is helpful to explore these relationships graphically, the hierarchical model allows us to determine whether these relationships are statistically significant, particularly in the presence of other competing summary series. We restrict our analysis of liquidity dynamics across multiple markets to the U.S. equity markets. We did this in part because of data consistency issues (there were no missing price-impact data for the U.S. equity markets over the period of interest) and because these markets exhibit somewhat consistent behavior, as seen in Figure 10. After visually exploring a number of candidates, we selected the 11 summary series described in Table 3.

We test the ability of the summary series to recover the liquidity dynamics across these markets in two ways. First we calculate a hit rate, which is the proportion of the time that the probit model, based solely on the summary series, accurately predicts the state identified by each of the underlying univariate models (or we count the proportion of time that we accurately predict the state of  $D_{it}$  for each  $i$  and  $t$  using the current estimate of  $\beta$ ,  $\Sigma$  and the summary series data  $x_{it}$ ). The naive hit rate is 33 percent, assuming random guessing, and the posterior average of probit model's hit rate was 66 percent indicating that the summary series are explaining a substantial portion of the liquidity dynamics. Second, we plotted the predicted probability of being in each state for each time point, using the probit model, against the average probability of being in each state for each time point, as shown in Figure 11. The way the predicted probabilities closely tracks the average probabilities confirm, again, the ability of the summary series to explain the liquidity dynamics.

We standardized the summary series (mean-centered and divided by the standard deviation), to compare the parameter estimates from the probit portion of the hierarchical model, seen in Table 4, directly with respect to their size. Because we force the latent

Table 3: Summary Series for the Hierarchical (Multiple Market) Model

Variable	Description
3-month Repo Rate	ICAP General Collateral Treasury 3-month repurchase agreement rate
Yield Curve	Yield on the constant maturity 10-year U.S. Treasury bond minus the yield on the constant maturity 2-year U.S. Treasury note
TED Spread	3-month LIBOR rate minus the 3-month U.S. Treasury bill yield
Moody's Baa Corporate Bond Index	Yield on the Moody's investment grade long-term corporate bond index
VIX <sup>®</sup> Index	Reflects the market estimate of future (30-day) volatility of the S&P 500
Dow Jones U.S. Real Estate Index	Index of real estate investment trusts (REITs) and other companies investing directly or indirectly in real estate through development, management, or ownership
S&P 500 Price-to-Book Ratio	Ratio of equity market value to book value per share of the S&P 500
Three-month LIBOR-OIS Spread	Difference between the 3-month LIBOR and the 3-month U.S. dollar overnight index swap (OIS) rate
5-year U.S. Breakeven Inflation Rate	Calculated by subtracting the real yield of the 5-year inflation-linked maturity curve from the yield of the closest 5-year nominal Treasury maturity. The result is the market-implied inflation expectation over the next 5 years
WTI Front-Month Price	Futures price for the near-dated expiry of the WTI oil contract
U.S. Dollar Index	Indicates the general international value of the U.S. dollar, by averaging exchange of the dollar against other major currencies
Source: OFR Analysis	

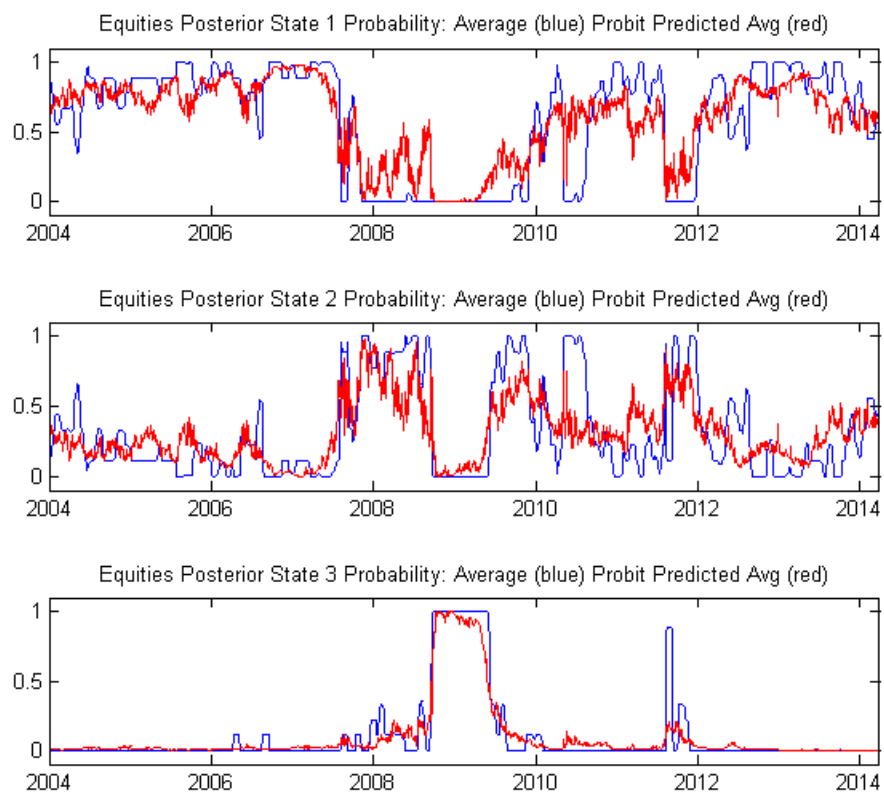


Figure 11: Average State Probabilities vs. Probit Predicted Probabilities  
 Sources: CRSP, WRDS, Bloomberg L.P, OFR analysis

value for state 1 to always be zero, to address the additive identification restriction, we only get parameter estimates for states 2 and 3 (which are really the difference between the unrestricted parameters of each of these states relative to state 1). The negative intercepts indicate that state 1, the low price-impact or high-liquidity state, is the most prevalent state when the associated summary variable is positive, and the fact that the intercept for state 3 is more negative than for state 2 indicates that state 3, the low-liquidity state, is the least likely state when the associated summary variable is positive.

Table 4: Posterior Parameter Estimates Probit Portion of Hierarchical Model

Summary Variable	Posterior Mean		Posterior StDev		t-Stat	
	State 2	State 3	State 2	State 3	State 2	State 3
Intercept	-0.64**	-1.01**	0.02	0.01	-40.27	-102.11
VIX <sup>®</sup>	0.62**	0.26**	0.03	0.02	23.77	15.67
WTI	0.83**	-0.23**	0.03	0.01	30.65	-16.53
3m Repo Rate	0.68**	-0.41**	0.02	0.01	38.28	-49.12
TED Spread	0.49**	-0.09**	0.03	0.01	18.53	-10.98
Yield Curve (10y–2y)	0.19**	-0.38**	0.02	0.04	9.66	-56.07
S&P 500 P/B Ratio	0.68**	-0.13**	0.03	0.01	23.85	-9.86
Dow Jones Real Estate Index	-1.17**	0.13**	0.02	0.01	-64.48	11.69
Moody's Baa Bond Index	-0.67**	0.47**	0.02	0.01	-43.57	60.83
LIBOR–OIS Spread	-0.64**	0.13**	0.05	0.02	-13.07	7.38
U.S. Dollar Index	-0.39**	-0.37**	0.04	0.01	-10.67	-41.05
U.S. 5y Breakeven Inflation	-0.03	0.00	0.02	0.01	-1.43	0.36

\*\* Significant at a 99 percent confidence level

Sources: CRSP, WRDS, Bloomberg L.P, OFR analysis

All but one of the summary series are statistically significant (only the U.S. five-year breakeven inflation is not significant in distinguishing the high-liquidity state from the mid-liquidity state). Within these results, there are some interesting patterns to note. First, there is a natural grouping among the summary series with regards to the pattern of the signs for the state 2 and 3 parameter estimates. As might be expected, VIX<sup>®</sup> has a positive-positive pattern indicating that higher levels of VIX<sup>®</sup> are associated with a higher probability of entering states with low liquidity. The persistence of high levels

of VIX<sup>®</sup> after the crisis makes VIX<sup>®</sup> more strongly related to the middle liquidity state as opposed to the crisis state. Another group of five summary series (WTI, three-month repo rate, S&P 500 P/B Ratio, TED Spread and Yield Curve (10-year minus 2-year yield)) exhibit a positive-negative pattern indicating that elevated levels of these summary series lead to a high probability of being in the middle liquidity state. Although this pattern may seem counter-intuitive, we observe that market movements, government actions and actions by central banks distorts and lowers these summary series during times of crisis, which corresponds to times of low liquidity (i.e., state 3). The next set of summary series (Dow Jones U.S. Real Estate Index, Moody's Baa Corporate Bond Index, and LIBOR-OIS spread) exhibit a negative-positive pattern which identifies them as measurements which have persistently high levels during times of low liquidity and then bounce back sufficiently during the moderate times of low liquidity to have the extremes associated with the crisis. Finally, the U.S. Dollar Index shows a significant negative-negative pattern, indicating that higher levels of the dollar are associated with higher probability of entering a high-liquidity state. This is consistent with a flight to quality, in which capital flows into the United States during episodes of stress, simultaneously pushing up the value of the dollar and flooding the domestic market with liquidity.

To assist in our understanding of these parameter estimates, we can compare the time-series plot for individual summary series versus the probit-predicted probability for each state. For example, Figure 12 presents this comparison for the TED spread. The spread remains low until mid-2007 and returns to persistent low levels in 2010. The early episode corresponds to consistently high probabilities for the high-liquidity state (i.e., state 1), consistent with the TED spread's role as a bellwether for funding liquidity. Between August 2007 and September 2008, when the TED spread begins to widen, the probability of state 2 jumps, supporting the positive coefficient in Table 4. After September 2008, the TED spread recedes relatively quickly from its peak, compared with the probability of



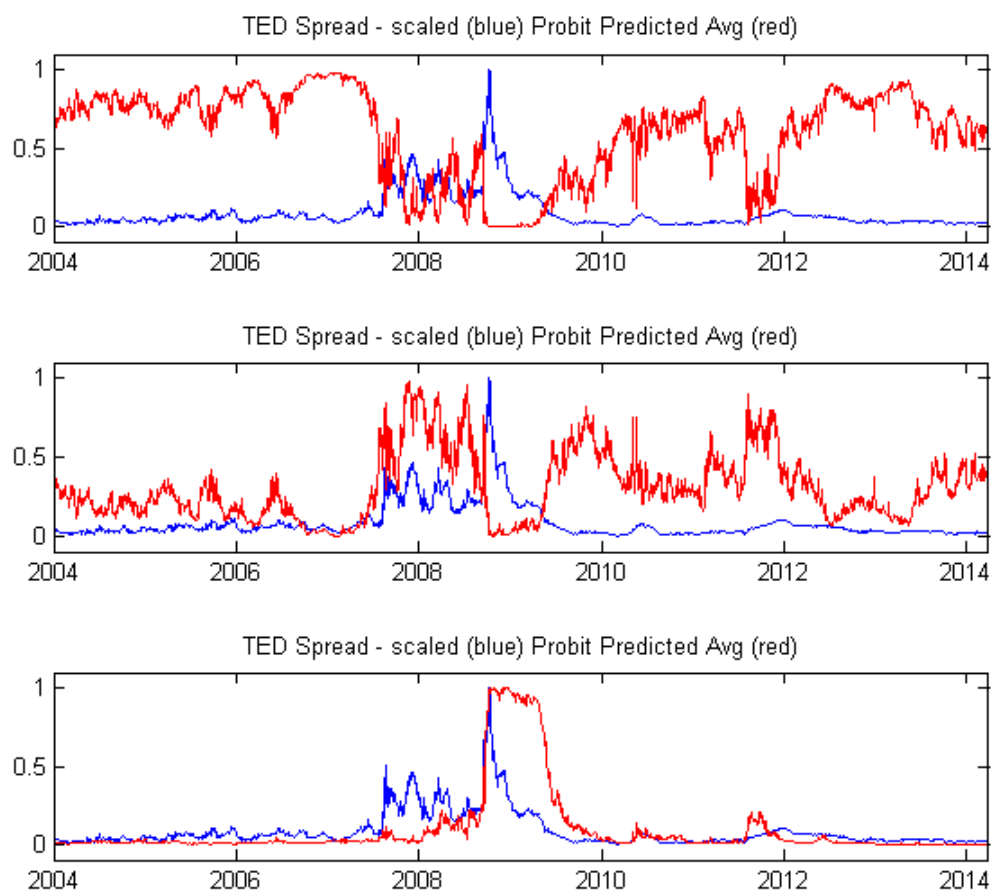


Figure 12: TED Spread and Average Probit Probabilities

Sources: CRSP, WRDS, Bloomberg L.P., OFR analysis

being in state 3, which remains elevated for the next year. This deviation is consistent with the negative coefficient on state 3 in Table 4. In contrast, the VIX<sup>®</sup> index has positive coefficients on both states 2 and 3 in Table 4. The VIX<sup>®</sup> is more persistently high after the 2008 shock, consistent with the positive coefficient on state 3 in Table 4. Moreover, it remains moderately elevated for much of the post-crisis period after 2009, when the probit-predicted probability for state 2 is also raised.

The next four (Moody's Baa Corporate Bond Index, VIX<sup>®</sup>, LIBOR and WTI) have a negative-positive pattern and are clear predictors of periods of low-liquidity. The VIX<sup>®</sup> has the strongest parameter estimate in absolute value and when the VIX<sup>®</sup> is high, the probability of being in state 2 drops and the probability of being in state 3 rises dramatically as indicated by plotting VIX<sup>®</sup> against the predicted probabilities in Figure 13.

The final summary variable (S&P 500 P/B ratio) has a positive-positive pattern (although only the parameter for state 3 is statistically significant). This suggests that, as the price of equities becomes large relative to the underlying book value of the firms, the system tends to be in state 3, the low-liquidity state. One possible explanation for this is that a high price-to-book ratio reflects a potential asset bubble, which could lead investors to engage in herd behavior (e.g., piling into different individual stocks and driving up returns then pulling out suddenly causing large price drops); this could cause not only increased volatility but also larger price-impacts.

Clearly, the hierarchical model provides interesting insights into how summary variables relate to liquidity dynamics and offers a valuable tool for further investigating and understanding the drivers of liquidity across a wide range of markets.

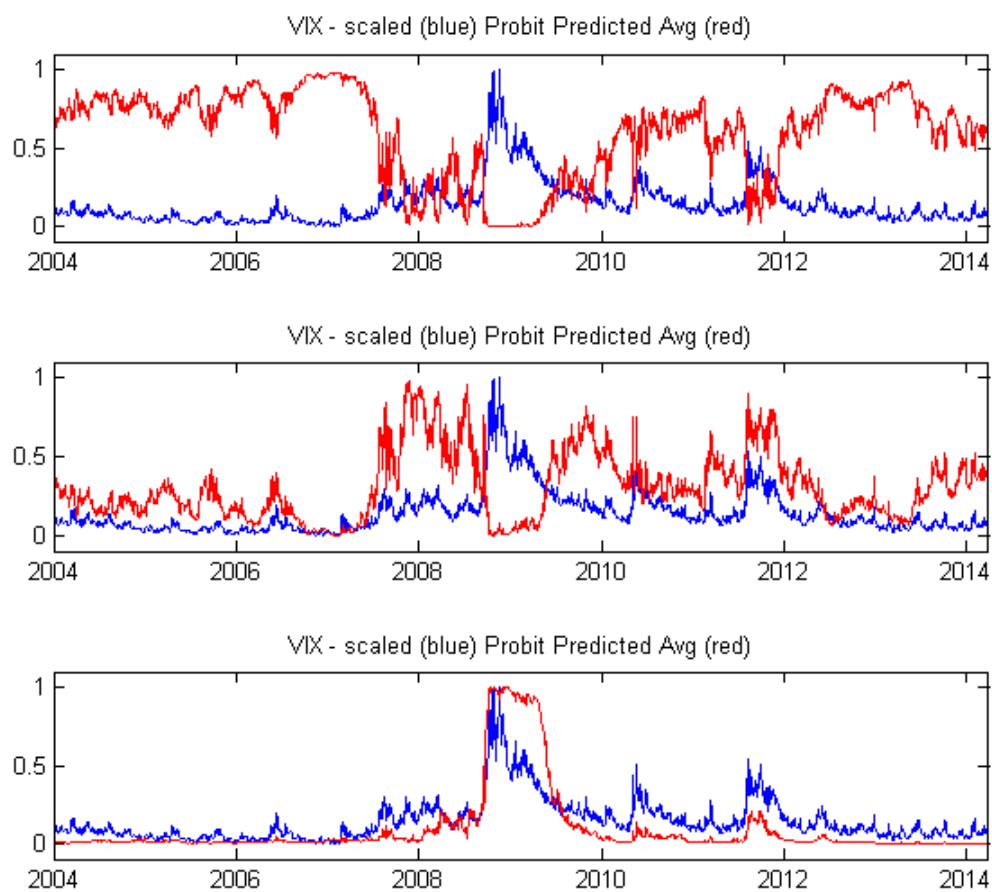


Figure 13: VIX<sup>®</sup> and Average Probit Probabilities  
 Sources: CRSP, WRDS, Bloomberg L.P., OFR analysis

## 5 Predicting Liquidity Regimes

An established, alternative approach to our proposed multivariate, hidden Markov model for modeling the price impacts across multiple markets is a vector auto-regressive (VAR) model, where the vector of price impacts across the various markets being considered are assumed to have an auto-correlation structure, meaning that lagged values of price impacts from one market potentially drives the current price impacts observed in some or all of the markets. Preliminary analysis of the same set of price impacts, across the equity markets, revealed that these price impacts are highly correlated. While only just over 70 percent of the variation is explained by the first factor, using a principal component analysis, Figure 14 illustrates that this factor captures the critical parts of the liquidity dynamics, with respect to the financial crisis.

This high level of multicollinearity in the data affects the Bayesian analysis of a VAR model (see Koop and Korobilis [2010] for a discussion of methods used), in that all of the lagged values (up to at least three lags) for each market are statistically significant for all of the markets.<sup>37</sup> While a VAR may offer some predictive power, it does not lead to a parsimonious representation of the liquidity dynamics across multiple markets, especially from the perspective of offering a useful tool for policy makers. If a linear model is to be used, a more parsimonious approach would be to model the primary latent factor, using a Bayesian factor model, and allow the transition dynamics of that latent factor to depend upon the summary series used in the multivariate, hidden Markov model.

Because of the difficulties providing policy makers with a parsimonious representation based a VAR or latent factor model, we choose to explore the predictive power of our proposed multivariate model and leave comparison of the relative predictive performance of these competing methods as a point for future research. The relatively high hit rate

---

<sup>37</sup>This was true even when we used the stochastic-search variable selection method proposed by George and McCulloch [1993], which achieves parsimony by removing variables with no relationships.

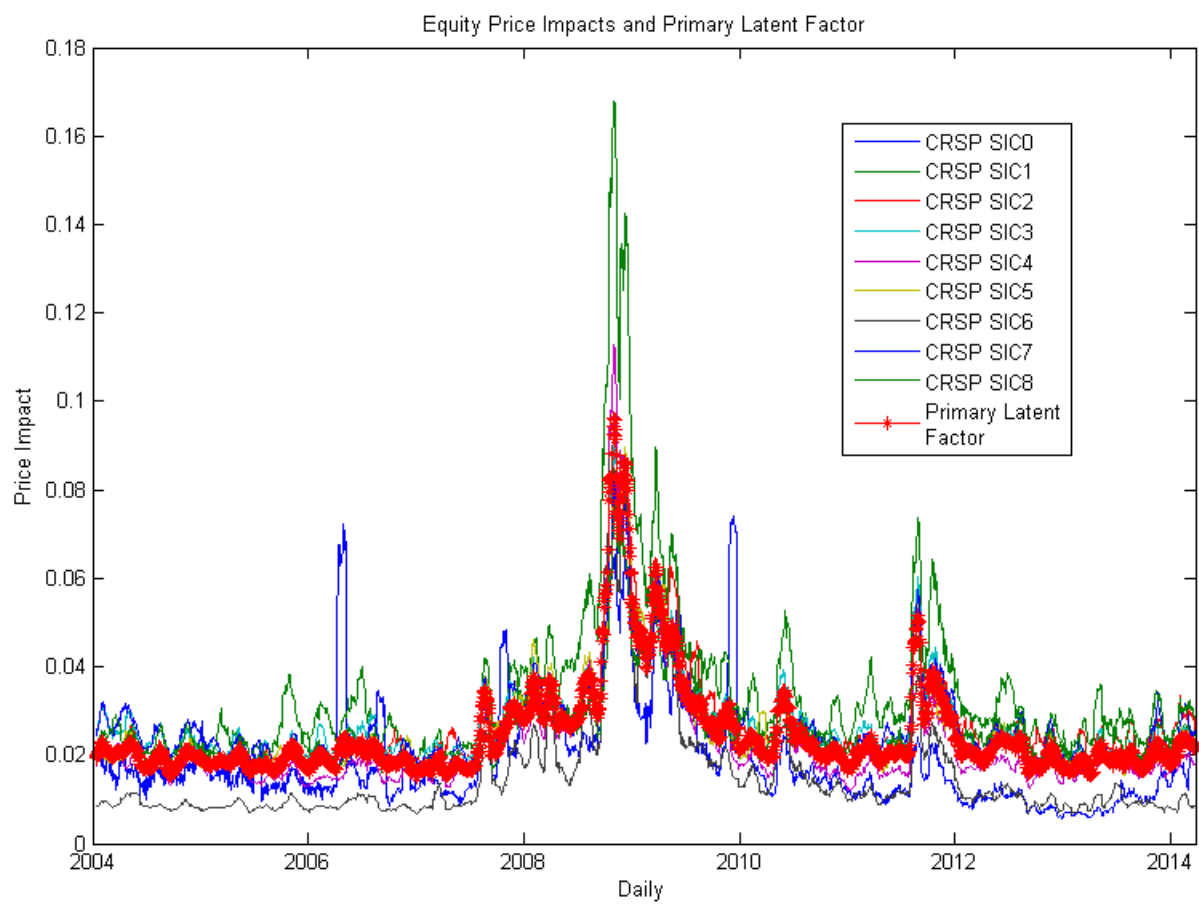


Figure 14: Price Impacts and Primary Liquidity Factor  
 Sources: CRSP, WRDS, OFR analysis

of the multivariate model suggests that it may have reasonable predictive power, and in fact the high hit rate, combined with the high temporal correlation in the summary series offers what appears to be a potentially powerful tool for predicting future periods of high price impact, which correspond to times of severe financial stress.

Table 5: Lagged Summary Series Hit Rates for Equity Price Impacts

		Predict CRSP Equity Price Impacts	
	Full Series	2004-05-01 to 2008-09-01 (Pre-Lehman)	2004-05-01 to 2008-02-21 (Pre-Bear Stearns)
Lag0	0.66	0.58	0.56
Lag1	0.65	0.58	0.56
Lag2	0.66	0.59	0.56
Lag3	0.65	0.59	0.56
Lag4	0.65	0.58	0.57
Lag5	0.65	0.59	0.57
Lag10	0.65	0.59	0.56
Lag15	0.65	0.58	0.56
Sources: CRSP, WRDS, Bloomberg L.P, OFR analysis			

There may be some concern that using the current value of the summary series to explain the current liquidity state may result in a model where any relationships that were found are endogenous or driven in part by common unobserved factors. In exploring this issue, we found that adding lagged values of the summary series, in addition to the current summary series, provide almost no additional benefit with regards to the ability to recover the average liquidity state. This led us to investigate using only lagged values of the summary series and excluding the current values of the summary series. Table 5, shows that using only lagged summary information that is up to five days old provides hit rates that are essentially the same as hit rates using current summary information. Even at lags of 10 and 15 days the model still recovers a substantial amount of the hit rate that occurs from using the current values of the summary series.

To illustrate the potential predictive power of the multivariate hidden Markov model during a period of financial stress, we calibrated the model by lagging the summary series at four different levels, using data from March 1, 2004 up through June 29, 2007: (i) the current values (lag 0); (ii) one day old (lag 1); (iii) five days old (lag 5); and (iv) 15 days old (lag 15). We then used the parameters estimated from each analysis and the prevailing summary series going forward to predict the probability of being in each of the different liquidity regimes. To clarify, using this approach, we input the July 1, 2007 summary data with the lag 0 analysis and predict the probabilities of being in the different liquidity regimes for July 1, 2007. On the other hand, using the same summary data with the lag 15 model we were predicting the probability of being in the different liquidity regime for July 16, 2007. As indicated by the dashed lines (which represent the predicted probabilities), in Figure 15, all the lagged models predict a jump to the high price impact or low liquidity regime well in advance of the crisis of 2008.

It is interesting to note that the Lag 5 and Lag 15 predictions identify the period of stress sooner. This happens in part because of the fact that they skip forward and in part because they are more sensitive to extreme jumps — perhaps because there is some endogeneity in the current summary series data and the current liquidity states, which is eliminated by lagging the summary series data at least five days.

Clearly these predictive results suggest that we explore ways to extend this predictive approach and contrast it with alternative predictive methods. Part of our extension will include overcoming data consistency issues with price impact measurements from markets beyond the equity markets. One approach would be to create a Bayesian Factor model, where the missing price impacts are treated as unobserved (latent) values that can be estimated in a manner that is consistent with the factor structure that is uncovered. This Bayesian factor model, can then be integrated into an extended version of the multivariate hidden Markov model, where the a finite mixture structure is used to group markets based

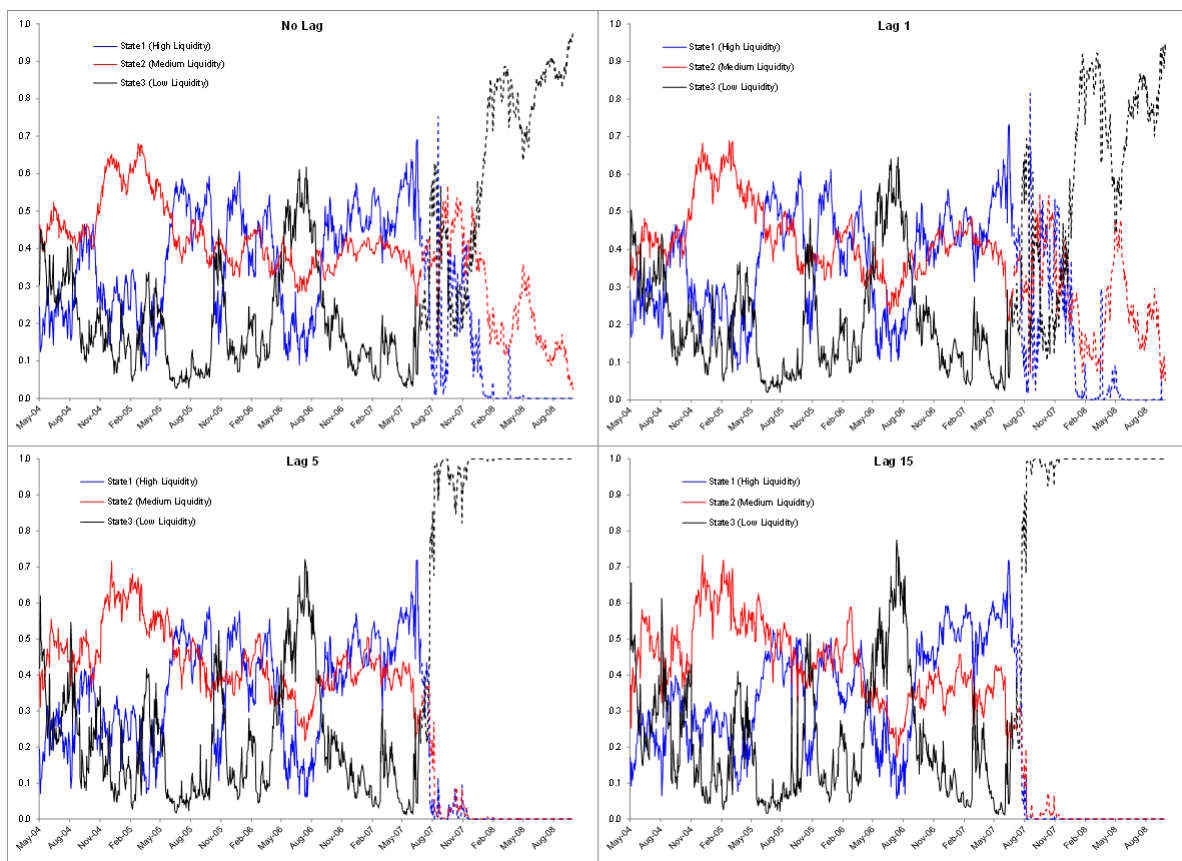


Figure 15: Predicted Liquidity States at Four Lags  
Sources: CRSP, WRDS, Bloomberg L.P., OFR analysis



on both their loadings on the underlying factors and the dynamics of their hidden Markov chains. Clearly it would be of interest to compare this extended model with the predictive performance of a version of the Bayesian factor model where the dynamics of the factors are driven by the summary series that are used with the multivariate hidden Markov model. Finally, to explore the practical value of this approach to policy makers, we need to explore the performance of this modeling for a range of different periods of extreme liquidity shocks both across time and across different geographic areas. We anticipate that this exploration will lead us to conclude that the impact of different summary series on predicting liquidity regimes change over time, ultimately leading us to develop a comprehensive model that has a dynamic component, with regards to the parameters in the multinomial probit portion of the multivariate hidden Markov model.

## 6 Conclusion

Liquidity is an elusive, yet essential component of the modern financial system. It is elusive because conceptually it is hard to define, and empirically it is hard to measure and predict. We attribute the challenges in liquidity measurement to three fundamental aspects of the phenomenon. Liquidity is latent, in the sense that the episodes of illiquidity we seek to understand are rare, and often emerge with little apparent warning. Liquidity is nonlinear, in the sense that price impact does not respond proportionately to additional order flow, making it difficult to extrapolate from ordinary markets to the behavior of those markets under stress. Liquidity is endogenous, in the sense that it often emerges as a positive externality in very active markets, making those busy venues attractive to others who seek the assurance that counterparties will be available when needed.

We address the challenges of latency, nonlinearity and endogeneity statistically with a Bayesian estimation of a hidden Markov chain individually for 33 separate time series

covering the CRSP and TRACE universes of U.S. equities and corporate bonds, plus multiple expiries of two key futures contracts, the VIX<sup>®</sup> volatility contract and the WTI oil contract. Three latent states (high, medium, and low price impact) are adequate to capture the observed liquidity structure of all 33 univariate series.

We also look for cross-sectional structure in the data by estimating a hierarchical Bayesian model, and testing the ability of several systemwide market summaries to recover the estimated aggregate liquidity dynamics. This exercise also permits an attribution of those estimated aggregate dynamics to meaningful economic interpretations. We explore the predictive power of this model, using lagged values of the summary series, and find that the model offers a possible predictive tool for identifying future jumps in market liquidity as far out as 15 days in advance. For reasons of data consistency, we have limited our initial efforts in this area to the U.S. equities markets.

Our results at this stage are preliminary, but also promising. In addition to testing for robustness and sensitivity, we see several immediate avenues for future research, including expanding the cross section of asset markets in the scope of analysis, comparing in more detail the liquidity behavior of wholesale funding markets, and experimenting with alternative portfolio formation rules.

## References

- Tobias Adrian and Hyun Song Shin. Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418–437, July 2010. URL <http://www.sciencedirect.com/science/article/pii/S1042957308000764>.
- Tobias Adrian, Paolo Colla, and Hyun Song Shin. Which financial frictions? parsing the evidence from the financial crisis of 2007-9. *NBER Macroeconomics Annual*, 27:159–214, May 2013a. URL <http://www.princeton.edu/~hsshin/www/whichfinancialfrictions.pdf>.
- Tobias Adrian, Daniel Covitz, and Nellie Liang. Financial stability monitoring. Finance and Economics Discussion Series 2013-21, Board of Governors of the Federal Reserve, 2013b. URL <http://www.federalreserve.gov/pubs/feds/2013/201321/201321pap.pdf>.
- James H Albert and Siddhartha Chib. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88(422):669–679, 1993. doi: 10.1080/01621459.1993.10476321. URL <http://www.tandfonline.com/doi/abs/10.1080/01621459.1993.10476321>.
- Franklin Allen and Douglas Gale. *Understanding Financial Crises*. Oxford University Press, 2009. URL <https://global.oup.com/academic/product/understanding-financial-crises-9780199251421?cc=us&lang=en&>.
- Y. Amihud. Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5:31–56, 2002. URL [https://umdrive.memphis.edu/pjain/Microstructure/Papers/Illiquidity%20and%20stock%20returns\\_Amihud.pdf](https://umdrive.memphis.edu/pjain/Microstructure/Papers/Illiquidity%20and%20stock%20returns_Amihud.pdf).
- Yakov Amihud and Haim Mendelson. Asset pricing and the bid-ask spread. *Journal of Financial Economics*, 17(2):223–249, 1986. doi: 10.1016/0304-405X(86)90065-6. URL <http://www.sciencedirect.com/science/article/pii/0304405X86900656>.
- Yakov Amihud, Haim Mendelson, and Lasse Heje Pedersen. *Market Liquidity: Asset Pricing, Risk, and Crises*. Cambridge University Press, 2013. URL <http://ebooks.cambridge.org/ebook.jsf?bid=CB09780511844393>.
- Adam Ashcraft, Morten L. Bech, and W. Scott Frame. The federal home loan bank system: The lender of next-to-last resort? *Journal of Money, Credit and Banking*, 42(4):551–583, June 2010. doi: 10.1111/j.

1538-4616.2010.00299.x. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1538-4616.2010.00299.x/abstract>.

Walter Bagehot. *Lombard Street: A Description of the Money Market*. H.S. King, 1873.

Bank of England. Instruments of macroprudential policy. Discussion paper, Bank of England, 2011. URL <http://www.bankofengland.co.uk/publications/other/financialstability/discussionpaper111220.pdf>.

Michael J. Barclay, Terrence Hendershott, and Kenneth Kotz. Automation versus intermediation: Evidence from treasuries going off the run. *Journal of Finance*, 61(5):2395–2414, October 2006. ISSN 00221082. URL <http://www.jstor.org/stable/3874714>.

Basel Committee on Banking Supervision. Basel iii: International framework for liquidity risk measurement, standards and monitoring. Technical Report 188, BIS, December 2010. URL <http://www.bis.org/publ/bcbs188.htm>.

Basel Committee on Banking Supervision. Basel iii: The liquidity coverage ratio and liquidity risk monitoring tools. Technical report, BIS, 2013. URL <http://www.bis.org/publ/bcbs238.htm>.

Leonard E. Baum, Ted Petrie, George Soules, and Norman Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *Annals of Mathematical Statistics*, 41(1):164–171, 1970. URL <http://projecteuclid.org/euclid.aoms/1177697196>.

Hendrik Bessembinder, William Maxwell, and Kumar Venkataraman. Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics*, 82(2):251–288, November 2006. URL <http://www.sciencedirect.com/science/article/pii/S0304405X06000699>.

Hendrik Bessembinder, Marios Panayides, and Kumar Venkataraman. Hidden liquidity: an analysis of order exposure strategies in electronic stock markets. *Journal of Financial Economics*, 94(3):361–383, December 2009. URL <http://www.sciencedirect.com/science/article/pii/S0304405X09001494#>.

Tim Bollerslev. Glossary to ARCH (GARCH). In *Volatility and Time Series Econometrics: Essays in Honor of Robert Engle*, Advanced Texts in Econometrics, chapter 8, pages 137–163. Oxford University Press, 2010. URL <https://global.oup.com/academic/product/volatility-and-time-series-econometrics-9780199549498?cc=us&lang=en&>.

- Richard Bookstaber, Michael D. Foley, and Brian F. Tivnan. Market liquidity and heterogeneity in the investor decision cycle. OFR Working Paper 15-02, Office of Financial Research, February 2015.
- Jean-Philippe Bouchaud, J. Doyne Farmer, and Fabrizio Lillo. How markets slowly digest changes in supply and demand. In T. Hens and K. Schenk-Hoppe, editors, *Handbook of Financial Markets: Dynamics and Evolution*, page 57156. Academic Press, 2008. URL <http://arxiv.org/pdf/0809.0822.pdf>.
- Kris Boudt, Ellen CS Paulus, and Dale WR Rosenthal. Funding liquidity, market liquidity and ted spread: A two-regime model. Working Paper 244, Free University of Brussels (VUB), November 2013. URL [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1668635](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1668635).
- Nicole M. Boyson, Christof W. Stahel, and Ren M. Stulz. Hedge fund contagion and liquidity shocks. *Journal of Finance*, 65(5):1789–1816, October 2010. URL <http://www.cob.ohio-state.edu/fin/faculty/stulz/publishedpapers/jofi1594.pdf>.
- Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of Markov Chain Monte Carlo*. CRC press, 2011. URL <http://www.crcpress.com/product/isbn/9781420079425>.
- M. Brunnermeier and L.H. Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, June 2009. URL <http://rfs.oxfordjournals.org/content/22/6/2201.abstract>.
- Markus K. Brunnermeier. Deciphering the liquidity and credit crunch 2007-2008. *Journal of Economic Perspectives*, 23(1):77–100, Winter 2009. URL <http://ideas.repec.org/a/aea/jecper/v23y2009i1p77-100.html>.
- Markus K. Brunnermeier and Yuliy Sannikov. Macroeconomics with a financial sector. Working Paper 236, National Bank of Belgium, October 2012. URL <file:///C:/Users/floodm/Downloads/SSRN-id2160894.pdf>.
- Charles Calomiris and Gary Gorton. The origins of banking panics: Models, facts, and bank regulation. In R. G. Hubbard, editor, *Financial Markets and Financial Crises*, chapter 4, pages 109–174. NBER and University of Chicago, 1991. URL <http://www.nber.org/chapters/c11484.pdf>.
- Olivier Cappé, Eric Moulines, and Tobias Rydén. *Inference in Hidden Markov Models*. Springer Verlag, 2005. URL <http://link.springer.com/book/10.1007/0-387-28982-8>.

- Chris K. Carter and Robert Kohn. On gibbs sampling for state space models. *Biometrika*, 81(3):541–553, September 1994. URL <http://biomet.oxfordjournals.org/content/81/3/541.short>.
- Han Chen, Vasco Cúrdia, and Andrea Ferrero. The macroeconomic effects of large-scale asset purchase programs. *Economic Journal*, 122(564):289–315, November 2012a. URL [http://www.newyorkfed.org/research/staff\\_reports/sr527.html](http://www.newyorkfed.org/research/staff_reports/sr527.html).
- Sally Chen, Philip Liu, Andrea Maechler, Chris Marsh, Sergejs Saksonovs, and Hyun Song Shin. Exploring the dynamics of global liquidity. Technical report, IMF, 2012b. URL <http://www.imf.org/external/pubs/ft/wp/2012/wp12246.pdf>.
- Tarun Chordia, Richard Roll, and Avanidhar Subrahmanyam. Commonality in liquidity. *Journal of Financial Economics*, 56(1):3–28, 2000. URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.139.1722&rep=rep1&type=pdf>.
- Marcia Millon Cornett, Jamie John McNutt, Philip E Strahan, and Hassan Tehranian. Liquidity risk management and credit supply in the financial crisis. *Journal of Financial Economics*, 101(2):297–312, August 2011. URL <http://www.sciencedirect.com/science/article/pii/S0304405X11000663>.
- Stefania D’Amico and Thomas B King. Flow and stock effects of large-scale treasury purchases: evidence on the importance of local supply. *Journal of Financial Economics*, 108(2):425–448, May 2012. URL <http://www.sciencedirect.com/science/article/pii/S0304405X12002425>.
- Stefania D’Amico, William English, David López-Salido, and Edward Nelson. The federal reserve’s large-scale asset purchase programs: Rationale and effects. *Economic Journal*, 122(564):415–46, November 2012.
- Tri Vi Dang, Gary Gorton, and Bengt Holmström. Financial crises and the optimality of debt for liquidity provision. Technical report, Yale U., May 2010. URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.183.7960&rep=rep1&type=pdf>.
- Tri Vi Dang, Gary Gorton, and Bengt Holmström. Ignorance, debt and financial crises. Technical report, Yale U., February 2012. URL [http://www.econ.cuhk.edu.hk/dept/chinese/seminar/12-13/1st-term/Paper\\_Ignorance.pdf](http://www.econ.cuhk.edu.hk/dept/chinese/seminar/12-13/1st-term/Paper_Ignorance.pdf).

- Piet De Jong and Neil Shephard. The simulation smoother for time series models. *Biometrika*, 82(2): 339–350, June 1995. URL <http://biomet.oxfordjournals.org/content/82/2/339.short>.
- Hans Degryse, Frank De Jong, and Vincent Van Kervel. The impact of dark trading and visible fragmentation on market quality. *Review of Finance*, page rfu027, 2014. URL <http://rof.oxfordjournals.org/content/early/2014/06/13/rof.rfu027.short>.
- J. Bradford DeLong, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance*, 45(2):379–395, June 1990. ISSN 00221082. URL <http://www.jstor.org/stable/2328662>.
- Jens Dick-Nielsen. Liquidity biases in trace. *Journal of Fixed Income*, 19(2):43–55, Fall 2009. doi: DOI: 10.3905/jfi.2009.19.2.043. URL <http://www.iijournals.com/doi/abs/10.3905/jfi.2009.19.2.043>.
- Feike C. Drost and Theo E. Nijman. Temporal aggregation of garch processes. *Econometrica*, 61(4): 909–927, July 1993. URL <http://www.jstor.org/stable/2951767>.
- Darrell Duffie. Presidential address: Asset price dynamics with slow-moving capital. *Journal of Finance*, 65(4):1237–1267, 2010. URL <http://www.darrellduffie.com/uploads/pubs/DuffieAFAPresidentialAddress2010.pdf>.
- D. Easley, N. Kiefer, M. OHara, and J. Paperman. Liquidity, information, and infrequently traded stocks. *Journal of Finance*, 51(4):1405–1436, September 1996. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.1996.tb04074.x/>.
- David Easley and Maureen O’Hara. Time and the process of security price adjustment. *Journal of Finance*, 47(2):577–605, June 1992. ISSN 00221082. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.1992.tb04402.x/abstract>.
- Nicholas Economides. The economics of networks. *International Journal of Industrial Organization*, 14(6):673–699, October 1996. URL <http://www.sciencedirect.com/science/article/pii/0167718796010156>.
- Douglas J. Elliott, Greg Feldberg, and Andreas Lehnert. The history of cyclical macroprudential policy in the united states. Working Paper 0008, OFR, May 2013. URL <http://www>.

treasury.gov/initiatives/ofr/research/Documents/OFRwp0008\_ElliottFeldbergLehnert\_AmericanCyclicalMacroprudentialPolicy.pdf.

Robert Engle, Michael Fleming, Eric Ghysels, and Giang Nguyen. Liquidity, volatility, and flights to safety in the us treasury market: Evidence from a new class of dynamic order book models. Staff Report 590, Federal Reserve Bank of New York, December 2012. URL [http://www.nyfedeconomists.org/research/staff\\_reports/sr590.pdf](http://www.nyfedeconomists.org/research/staff_reports/sr590.pdf).

Robert F. Engle. The econometrics of ultra-high-frequency data. *Econometrica*, 68(1):1–22, January 2000. ISSN 00129682. URL <http://www.jstor.org/stable/2999473>.

Robert F Engle and Joe Lange. Predicting vnet: A model of the dynamics of market depth. *Journal of Financial Markets*, 4(2):113–142, April 2001. URL <http://www.sciencedirect.com/science/article/pii/S1386418100000197>.

Andrew J. Entwistle and Jonathan H. Beemer. Approaches to asset recovery for pension fund subprime exposure. Technical report, Entwistle & Cappucci, February 2008. URL <http://www.entwistle-law.com/news/publications/000041>.

Martin D.D. Evans and Richard K. Lyons. Time-varying liquidity in foreign exchange. *Journal of Monetary Economics*, 49(5):1025–1051, 2002. ISSN 0304-3932. doi: [http://dx.doi.org/10.1016/S0304-3932\(02\)00124-1](http://dx.doi.org/10.1016/S0304-3932(02)00124-1). URL <http://www.sciencedirect.com/science/article/pii/S0304393202001241>.

Emmanuel Farhi and Jean Tirole. Leverage and the central banker’s put. *American Economic Review*, 99(2):589–593, May 2009. ISSN 00028282. URL <http://www.jstor.org/stable/25592463>. Papers and Proceedings of the One Hundred Twenty-First Meeting of the American Economic Association.

FINRA. Historic trace data: Enhanced historic time and sales trade record file layout. Technical report, FINRA, 2009. URL <http://www.finra.org/web/groups/industry/@ip/@comp/@mt/documents/appsupportdocs/p120989.pdf>.

Michael J. Fleming. Federal reserve liquidity provision during the financial crisis of 2007-2009. *Annual Review of Financial Economics*, 4:161–177, 2012. doi: 10.1146/annurev-financial-110311-101735. URL <http://www.annualreviews.org/doi/abs/10.1146/annurev-financial-110311-101735>.



- Sean Foley, Katya Malinova, and Andreas Park. Dark trading on public exchanges. Technical report, University of Toronto, February 2013. URL <http://ssrn.com/abstract=2182839>.
- Thierry Foucault, Marco Pagano, and Ailsa Röell. *Market Liquidity: Theory, Evidence, and Policy*. Oxford University Press, 2013. URL <http://ukcatalogue.oup.com/product/9780199936243.do>.
- Xavier Gabaix, Parameswaran Gopikrishnan, Vasiliki Plerou, and H Eugene Stanley. Institutional investors and stock market volatility. *Quarterly Journal of Economics*, 121(2):461–504, 2006. URL <http://qje.oxfordjournals.org/content/121/2/461.abstract>.
- Alexandros Gabrielsen, Massimiliano Marzo, and Paolo Zagaglia. Measuring market liquidity: An introductory survey. Working Paper DSE 802, U. di Bologna, 2011. URL <http://arxiv.org/pdf/1112.6169.pdf>.
- John Geanakoplos. Liquidity, default, and crashes: Endogenous contracts in general equilibrium. In Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, editors, *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume II*, 2003.
- Deborah Gefang, Gary Koop, and Simon M. Potter. Understanding liquidity and credit risks in the financial crisis. *Journal of Empirical Finance*, 18(5):903–914, December 2011. URL <http://www.sciencedirect.com/science/article/pii/S0927539811000508>.
- Andrew Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin. *Bayesian Data Analysis*. Chapman & Hall/CRC, 3rd edition, 2013. URL <http://www.amazon.com/dp/1439840954/>.
- Edward I George and Robert E McCulloch. Variable selection via gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889, 1993. doi: 10.1080/01621459.1993.10476353. URL <http://amstat.tandfonline.com/doi/abs/10.1080/01621459.1993.10476353#.VMorSf1dVCY>.
- Charles AE Goodhart. The background to the 2007 financial crisis. *International Economics and Economic Policy*, 4(4):331–346, February 2008. URL <http://link.springer.com/article/10.1007%2Fs10368-007-0098-0?LI=true>.
- Joel Hasbrouck. Intraday price formation in us equity index markets. *Journal of Finance*, 58(6):2375–2400, December 2003. URL <http://onlinelibrary.wiley.com/doi/10.1046/j.1540-6261.2003.00609.x/abstract>.

- Joel Hasbrouck. *Empirical Market Microstructure: The Institutions, Economics, and Econometrics of Securities Trading*. Oxford University Press, 2007. URL <https://global.oup.com/academic/product/empirical-market-microstructure-9780195301649>.
- Joel Hasbrouck and Gideon Saar. Technology and liquidity provision: The blurring of traditional definitions. *Journal of Financial Markets*, 12(2):143–172, 2009. ISSN 1386-4181. doi: <http://dx.doi.org/10.1016/j.finmar.2008.06.002>. URL <http://www.sciencedirect.com/science/article/pii/S1386418108000220>.
- Joel Hasbrouck and Duane J Seppi. Common factors in prices, order flows, and liquidity. *Journal of Financial Economics*, 59(3):383–411, 2001. URL <http://www.sciencedirect.com/science/article/pii/S0304405X0000091X>.
- Craig W. Holden, Stacey E. Jacobsen, and Avanidhar Subrahmanyam. The empirical analysis of liquidity. *Foundations and Trends in Finance*, forthcoming, 2014. doi: 10.2139/ssrn.2402215. URL [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2402215](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2402215).
- Bengt Holmström. The nature of liquidity provision: When ignorance is bliss. Presidential address, Econometric Society, 2012. URL <http://economics.mit.edu/files/7500>.
- Bengt Holmström and Jean Tirole. *Inside and Outside Liquidity*. MIT Press, 2013. URL <http://www.amazon.com/dp/0262518538/>.
- Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1):65–91, March 1993. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.1993.tb04702.x/full>.
- Narasimhan Jegadeesh and Sheridan Titman. Profitability of momentum strategies: An evaluation of alternative explanations. *Journal of Finance*, 56(2):699–720, April 2001. URL <http://onlinelibrary.wiley.com/doi/10.1111/0022-1082.00342/abstract>.
- Rudolf E. Kálmán. A new approach to linear filtering and prediction problems. *Journal of Fluids Engineering*, 82(1):35–45, March 1960. doi: 10.1115/1.3662552. URL <http://fluidsengineering.asmedigitalcollection.asme.org/article.aspx?articleid=1430402>.

- G. Andrew Karolyi, Kuan-Hui Lee, and Mathijs A. van Dijk. Understanding commonality in liquidity around the world. *Journal of Financial Economics*, 105:82–112, 2012. URL <http://www.sciencedirect.com/science/article/pii/S0304405X11002844>.
- Ira G. Kawaller, Paul D. Koch, and Timothy W. Koch. The temporal price relationship between s&p 500 futures and the s&p 500 index. *Journal of Finance*, 42(5):1309–1329, December 1987. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.1987.tb04368.x/abstract>.
- Charles P. Kindleberger. *A Financial History of Western Europe*. Oxford University Press, 2 edition, 1993.
- Gary Koop and Dimitris Korobilis. *Bayesian multivariate time series methods for empirical macroeconomics*. Now Publishers Inc., 2010. URL <https://books.google.com/books?id=JcVn9IJ9wSsC>.
- A. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, November 1985. URL <https://www.econometricsociety.org/publications/econometrica/1985/11/01/continuous-auctions-and-insider-trading>.
- Albert S. Kyle and Anna A. Obizhaeva. Market microstructure invariants: Empirical evidence from portfolio transitions. working paper, University of Maryland, December 2011a. URL [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1978943](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1978943).
- Albert S. Kyle and Anna A. Obizhaeva. Market microstructure invariants: Theory and implications of calibration. working paper, University of Maryland, December 2011b. URL [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1978932](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1978932).
- Albert S. Kyle and Anna A. Obizhaeva. Large bets and stock market crashes. working paper, University of Maryland, May 2013. URL [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2023776](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2023776).
- Albert S. Kyle and Anna A. Obizhaeva. Market microstructure invariants: Theory and empirical tests. Working paper, University of Maryland, October 2014. URL [http://www.rhsmith.umd.edu/faculty/obizhaeva/Kyle\\_Obizhaeva\\_20130607.pdf](http://www.rhsmith.umd.edu/faculty/obizhaeva/Kyle_Obizhaeva_20130607.pdf).
- Sriketan Mahanti, Amrut Nashikkar, Marti Subrahmanyam, George Chacko, and Gaurav Mallik. Latent liquidity: A new measure of liquidity, with an application to corporate bonds. *Journal of Financial*

- Economics*, 88(2):272–298, May 2008. ISSN 0304-405X. doi: <http://dx.doi.org/10.1016/j.jfineco.2007.02.006>. URL <http://www.sciencedirect.com/science/article/pii/S0304405X08000172>.
- Miriam Marra. Illiquidity commonality across equity and credit markets. working paper, University of Reading, April 2013. URL [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2019512](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2019512).
- Robert McCauley. Benchmark tipping in the money and bond markets. *BIS Quarterly Review*, pages 39–45, March 2001. URL [http://www.bis.org/publ/r\\_qt0103.htm](http://www.bis.org/publ/r_qt0103.htm).
- Robert McCulloch and Peter E. Rossi. An exact likelihood analysis of the multinomial probit model. *Journal of Econometrics*, 64(1):207–240, September–October 1994. URL <http://www.sciencedirect.com/science/article/pii/0304407694900647>.
- Robert E. McCulloch, Nicholas G. Polson, and Peter E. Rossi. A bayesian analysis of the multinomial probit model with fully identified parameters. *Journal of Econometrics*, 99(1):173 – 193, 2000. ISSN 0304-4076. doi: [http://dx.doi.org/10.1016/S0304-4076\(00\)00034-8](http://dx.doi.org/10.1016/S0304-4076(00)00034-8). URL <http://www.sciencedirect.com/science/article/pii/S0304407600000348>.
- Perry Mehrling. *The New Lombard Street: How the Fed Became the Dealer of Last Resort*. Princeton University Press, November 2010. URL <http://www.amazon.com/dp/0691143986>.
- Francois-Louis Michaud and Christian Upper. What drives interbank rates? evidence from the libor panel. *BIS Quarterly Review*, pages 47–58, March 2008. URL [http://www.bis.org/publ/qtrpdf/r\\_qt0803f.htm](http://www.bis.org/publ/qtrpdf/r_qt0803f.htm).
- Stephen Morris and Hyun Song Shin. Liquidity black holes. *Review of Finance*, 8(1):1–18, 2004. URL <http://cowles.econ.yale.edu/P/cp/p11a/p1114.pdf>.
- Harold G. Moulton. Commercial banking and capital formation. *Journal of Political Economy*, 26(5, 6, 7, 9):484–508, 638–663, 705–731, 849–881, 1918.
- Ulrich A. Müller, Michel M. Dacorogna, Richard B. Olsen, Olivier V. Pictet, Matthias Schwarz, and Claude Morgenegg. Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis. *Journal of Banking and Finance*, 14(6):1189–1208, 1990. ISSN 0378-4266. doi: [http://dx.doi.org/10.1016/0378-4266\(90\)90009-Q](http://dx.doi.org/10.1016/0378-4266(90)90009-Q). URL <http://www.sciencedirect.com/science/article/pii/037842669090009Q>.

- Mahendrarajah Nimalendran and Sugata Ray. Informational linkages between dark and lit trading venues. *Journal of Financial Markets*, 17:230–261, January 2014. ISSN 1386-4181. doi: 10.1016/j.finmar.2013.02.003. URL <http://www.sciencedirect.com/science/article/pii/S1386418113000128>.
- NYU Volatility Institute (V-Lab). V-lab: Models. Technical report, New York University, 2014. URL <http://vlab.stern.nyu.edu/doc?topic=mdls>. Accessed 29 December 2014.
- Marco Pagano. Trading volume and asset liquidity. *Quarterly Journal of Economics*, 104(2):255–274, May 1989. URL <http://qje.oxfordjournals.org/content/104/2/255.short>.
- Christine A. Parlour and Duane J. Seppi. Limit order markets: A survey. In A.W.A. Boot and A.V. Thakor, editors, *Handbook of Financial Intermediation & Banking*, volume 5. Elsevier, 2008. URL [http://www.edegan.com/pdfs/Parlour%20Seppi%20\(2008\)%20-%20Limit%20Order%20Markets%20A%20Survey.pdf](http://www.edegan.com/pdfs/Parlour%20Seppi%20(2008)%20-%20Limit%20Order%20Markets%20A%20Survey.pdf).
- Carmen M. Reinhart and Kenneth Rogoff. *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press, September 2009. URL <http://press.princeton.edu/titles/8973.html>.
- Richard Roll. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39(4):1127–1139, September 1984. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.1984.tb03897.x/full>.
- Moritz Schularick and Alan M. Taylor. Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870–2008. *American Economic Review*, 102(2):1029–1061, April 2012. ISSN 00028282. URL <http://www.jstor.org/stable/23245443>.
- Hans R Stoll. Friction. *Journal of Finance*, 55(4):1479–1514, August 2000. URL <http://onlinelibrary.wiley.com/doi/10.1111/0022-1082.00259/abstract>.
- Suresh Sundaresan and Zhenyu Wang. Y2k options and the liquidity premium in treasury bond markets. Staff Report 266, Federal Reserve Bank of New York, November 2006. URL [http://www.fednewyork.org/research/staff\\_reports/sr266.pdf](http://www.fednewyork.org/research/staff_reports/sr266.pdf).
- J. B. Taylor and J. C. Williams. A black swan in the money market. *American Economic Journal: Macroeconomics*, 1:58–83, 2009.

Jean Tirole. Illiquidity and all its friends. *Journal of Economic Literature*, 49(2):287–325, 2011. URL [http://idei.fr/doc/by/tirole/illiquidity-may12\\_11.pdf](http://idei.fr/doc/by/tirole/illiquidity-may12_11.pdf).

Bence Toth, Yves Lempriere, Cyril Deremble, Joachim De Lataillade, Julien Kockelkoren, and J-P Bouchaud. Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2):021006/1–11, 2011. URL <http://prx.aps.org/pdf/PRX/v1/i2/e021006>.

Haoxiang Zhu. Do dark pools harm price discovery? *Review of Financial Studies*, 27(3):747–789, March 2013. doi: 10.1093/rfs/hht078. URL <http://dspace.mit.edu/openaccess-disseminate/1721.1/88124>.